

Kara, Engin

Working Paper

The Curse of Flexibility Under Uncertainty

CESifo Working Paper, No. 12166

Provided in Cooperation with:

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Kara, Engin (2025) : The Curse of Flexibility Under Uncertainty, CESifo Working Paper, No. 12166, Munich Society for the Promotion of Economic Research - CESifo GmbH, Munich

This Version is available at:

<https://hdl.handle.net/10419/331631>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

CES ifo

**12166
2025**

September 2025

Working Papers

The Curse of Flexibility Under Uncertainty

Engin Kara

CES ifo

Imprint:

CESifo Working Papers

ISSN 2364-1428 (digital)

Publisher and distributor: Munich Society for the Promotion
of Economic Research - CESifo GmbH

Poschingerstr. 5, 81679 Munich, Germany
Telephone +49 (0)89 2180-2740

Email office@cesifo.de
<https://www.cesifo.org>

Editor: Clemens Fuest

An electronic version of the paper may be downloaded free of charge

- from the CESifo website: www.ifo.de/en/cesifo/publications/cesifo-working-papers
- from the SSRN website: www.ssrn.com/index.cfm/en/cesifo/
- from the RePEc website: <https://ideas.repec.org/s/ces/ceswps.html>

The Curse of Flexibility Under Uncertainty

Engin Kara*

Abstract

How firms respond to uncertainty determines economic policy effectiveness. Using Brexit as a natural experiment, I document that flexible price-setters—those most responsive to monetary policy—paradoxically reduce adjustment more than sticky firms under uncertainty. This ‘curse of flexibility’ reverses menu cost models’ foundational prediction that flexibility amplifies responses. Under uncertainty, resetting prices exposes firms to symmetric shocks, while maintaining current prices provides partial insulation. Flexible-price firms can afford to exploit this differential exposure by waiting; sticky-price firms cannot. This creates a policy challenge: uncertainty weakens monetary transmission when needed most, as flexible firms—the most responsive channel—become more cautious during crises.

Keywords: Menu costs, price stickiness, uncertainty, state-dependent pricing, heterogeneous firms, monetary policy transmission.

JEL classification numbers: E31, D83, E52, D21, L11.

*Corresponding author: Cardiff Business School, University of Cardiff, Cardiff, CF10 3EU, UK. E-mail: karae1@cardiff.ac.uk

1. Introduction

Economic theory holds that flexibility facilitates adjustment. In standard menu cost models, firms with lower adjustment costs respond more readily to shocks. This principle, established in seminal work by Sheshinski and Weiss (1977) and Caplin and Spulber (1987), shows that when volatility rises, flexible firms should become even more responsive and their ability to respond quickly becomes increasingly valuable. This relationship underpins our understanding of price stickiness and monetary transmission: reducing frictions enables responsiveness, and this relationship strengthens during turbulent times. Indeed, modern state-dependent models from Golosov and Lucas (2007) to Alvarez et al. (2016) build on this foundation, predicting that lower menu costs amplify responses to aggregate shocks. Generating state-dependent price adjustment—where firms respond differently to shocks depending on aggregate conditions—has become a litmus test for models: time-dependent models like Calvo (1983) are often challenged precisely because adjustment probabilities remain constant regardless of economic conditions.

This paper documents a fundamental reversal of this relationship. Using comprehensive UK micro-price data (2003–2021) covering three distinct uncertainty episodes (Financial Crisis, Brexit, COVID-19), I show that flexible firms reduce their adjustment probability significantly more than sticky firms when uncertainty rises—the opposite of theoretical predictions. For the full sample, flexible firms reduce adjustment by approximately 1 percentage point more than sticky firms when uncertainty rises by one standard deviation—a differential comparable in magnitude to monetary policy shocks. During Brexit, this differential nearly doubles to 1.8 percentage points, suggesting the curse intensifies with the severity of uncertainty shocks. Remarkably, the curse completely reverses under monetary policy shocks: the same flexible firms show a 0.91 percentage point adjustment advantage over sticky firms, confirming that the curse is state-dependent rather than reflecting fixed firm characteristics. The curse emerges systematically whenever uncertainty exceeds a critical threshold, transforming flexibility from an advantage into a liability.

My methodological contribution centres on exploiting heterogeneity in firm-level adjustment frequencies to identify differential responses to uncertainty. While previous work examines average effects, I test whether uncertainty affects firms differently based on their

flexibility type—their historical frequency of price adjustment. To address endogeneity, I employ predetermined flexibility measures calculated either before major uncertainty episodes (pre-2005) or during calm periods. Brexit provides particularly clean identification: firms faced elevated uncertainty about future trade terms while maintaining normal operations, allowing separation of strategic price-setting responses from operational constraints. These quasi-experimental approaches yield even stronger curse effects than the baseline, confirming that results reflect inherent firm characteristics, not shock-induced responses.

I develop a heterogeneous menu cost model to understand this mechanism. When uncertainty is low, firms near their adjustment thresholds are most sensitive to shocks—small changes push them over the edge. But high uncertainty reverses this intuition. Firms that reset prices today lock in a new position that will be equally vulnerable to upward and downward shocks. In contrast, firms maintaining existing prices have already drifted toward one adjustment boundary, making them more likely to hit that nearby threshold than the distant one—effectively facing risk from only one direction. This asymmetric exposure provides partial insulation. The difference matters most for flexible firms: with low adjustment costs, they can afford to maintain this advantageous position. Sticky firms, facing high costs, must adjust when price gaps become too large regardless. The result contradicts standard predictions: precisely when theory says flexible firms should leverage their advantage, they instead become the most cautious.

These findings reveal that uncertainty’s effect on price adjustment is far more complex than previously understood. While aggregate adjustment frequencies may appear stable (Gagliardone et al. (2025)) or even increase during uncertainty episodes (as in Vavra (2014)), this masks dramatic heterogeneous responses. The apparent aggregate stability results from compositional effects: flexible firms (typically 30% of firms) show large responses while sticky firms (70% of firms) barely change.

The curse of flexibility extends Bloom (2009)’s ‘wait-and-see’ dynamics to price-setting with a crucial difference: the effect is strongest precisely for firms with the lowest adjustment costs, reversing standard predictions. This helps reconcile puzzles in the literature. Tenreiro and Thwaites (2016)’s finding that monetary policy is less powerful in recessions may partly reflect flexible firms’ strategic inaction—policy loses its most effective transmission channel

when firms best able to adjust choose not to. The curse likely interacts with other frictions. Financial constraints (Gilchrist et al. (2017)) may force adjustment despite the curse, while digitalisation reduces menu costs economy-wide, potentially amplifying these effects. Understanding such interactions becomes increasingly important for policy design.

The remainder of the paper is organised as follows. Section 2 presents the data and measurement approach. Section 3 documents the curse of flexibility empirically, showing heterogeneous responses to uncertainty across firm types. Section 4 examines Brexit as a natural experiment, providing clean identification of the curse. Section 5 contrasts uncertainty with monetary policy shocks, demonstrating that the curse is state-dependent rather than a fixed firm characteristic. Section 6 develops a heterogeneous menu cost model explaining why conventional theory predicts the opposite of what I find. Section 7 analyses the differential exposure mechanism that generates the curse. Section 8 concludes with implications for monetary policy effectiveness during crises.

2. Empirical Evidence

2.1. Data and Measurement

I use the UK CPI micro-price dataset compiled by Davies (2021), covering from 1988 to 2021. For the analysis, I focus on 2003-2021 because expenditure weights are more reliable in this period (standard deviation of expenditure weights falls from 2.8 to 1.6).

Sample construction begins with aggregating monthly data to quarterly frequency, calculating log price changes between consecutive observations, expressed in percentage points. To mitigate the impact of extreme observations, I winsorise price changes at the 1st and 99th percentiles. Importantly, I exclude temporary sales to focus on regular pricing decisions.¹ For the baseline period (2003-2021), quarterly data coverage remains stable with approximately 90,000-115,000 price quotes per quarter. The baseline regression sample contains 6.3 million price change observations across 339,535 unique products, where each product represents a specific item variety (e.g., Royal Gala apples) at a particular store location.

¹The correlation between my inflation measure and official CPI is 0.6 when excluding sales, versus 0.4 when including them. The higher correlation without sales suggests that temporary price reductions add noise unrelated to underlying inflation dynamics, supporting my focus on regular pricing decisions.

The dependent variable is the non-sale price change ($\Delta p_{i,t}$), which excludes temporary price reductions. All price changes are computed as log differences, ensuring comparability across products with different price levels. I construct two key inflation measures:

$$\pi_t = \sum_i \left(\Delta p_{i,t} \times \frac{w_{i,t}}{\sum_j w_{j,t}} \right) \quad (1)$$

$$\pi_{c,t} = \sum_{i \in c} \left(\Delta p_{i,t} \times \frac{w_{i,t}}{\sum_{j \in c} w_{j,t}} \right) \quad (2)$$

where $w_{i,t}$ represents the expenditure weight of product i at time t , and c denotes the COICOP category. COICOP categories include Food, Alcohol, Clothing, Housing, Furnishing, Health, Transport, Communication, Recreation, Restaurants, and Miscellaneous goods and services.

2.1.1. Key Variables

Following Vavra (2014), I use cross-sectional price dispersion as my primary measure. Specifically, I measure price uncertainty using the weighted variance of price changes:

$$\sigma_t^2 = \sum_i w_{i,t} (\Delta p_{i,t} - \pi_t)^2 \quad (3)$$

While not a direct measure of fundamental uncertainty, dispersion captures the heterogeneous environment firms navigate when setting prices. When firms cannot directly observe competitors' costs or demand conditions, the distribution of observed prices provides their best signal of environmental volatility.

To ensure that uncertainty effects are not confounded by asymmetric risks, I control for the third moment of the price change distribution:

$$\text{Skew}_t = \frac{\sum_i w_{i,t} (\Delta p_{i,t} - \pi_t)^3}{(\sum_i w_{i,t} (\Delta p_{i,t} - \pi_t)^2)^{3/2}} \quad (4)$$

I use lagged values of these measures, standardised to have zero mean and unit standard

deviation:

$$\tilde{\sigma}_{t-1} = \frac{\sigma_{t-1}^2 - \overline{\sigma^2}}{\text{SD}(\sigma^2)} \quad (5)$$

$$\widetilde{\text{Skew}}_{t-1} = \frac{\text{Skew}_{t-1} - \overline{\text{Skew}}}{\text{SD}(\text{Skew})} \quad (6)$$

The use of lagged values serves two purposes. First, it addresses potential simultaneity concerns: current period price adjustments mechanically affect the contemporaneous distribution moments, but firms cannot influence past aggregate uncertainty when making current decisions. Second, it captures the realistic decision-making timeline where firms observe aggregate conditions and then decide whether to adjust prices in the following period. This standardisation facilitates interpretation: a one-unit change represents a one standard deviation change in the respective measure.

To test heterogeneous responses to uncertainty, I construct a unit-specific measure of price adjustment frequency, freq_i , defined as the ratio of price changes to total observations for each product-store combination. This measure provides the cross-sectional variation needed to identify the curse of flexibility.

2.1.2. Identification

My identification assumption is that high realised dispersion signals continued volatility to firms. When price changes become more dispersed, firms interpret this as evidence of an uncertain environment that will likely persist. This assumption follows standard practice in the uncertainty literature where realised volatility proxies for expected future volatility (Bloom (2009)).²

Price dispersion exhibits precisely the persistence structure needed for forward-looking price-setting. While the first-order autocorrelation is 0.465, the AR(2) specification reveals richer dynamics:

$$\tilde{\sigma}_t^2 = 0.185\tilde{\sigma}_{t-1}^2 + 0.528\tilde{\sigma}_{t-2}^2 + \varepsilon_t \quad (7)$$

The dominant second lag (0.528) indicates that dispersion follows multi-quarter episodes

²Unlike choice variables such as employment, price dispersion reflects market-wide conditions beyond any individual firm's control. While each firm contributes to dispersion through its pricing decision, no firm can manipulate the aggregate measure.

rather than simple decay. This persistence structure has direct implications for my mechanism: when firms observe high dispersion today, they can rationally infer they are entering an episode of elevated volatility that will persist for multiple quarters.³

To address endogeneity concerns, I employ three complementary identification strategies that leverage different sources of variation. First, I use lagged dispersion (σ_{t-1}) to ensure firms respond to pre-existing conditions rather than concurrent shocks. Second, as a robustness check, I recalculate adjustment frequencies using pre-2005 data or calm periods. Third, I exploit the Brexit referendum as a quasi-experimental shock, where the surprise outcome generated variation orthogonal to firms’ prior pricing decisions. During Brexit, dispersion increased from 0.086 to approximately 0.48-0.50—a nearly six-fold increase representing a shift of 0.4 standard deviations. This unexpected, sustained increase in dispersion—unrelated to operational constraints—provides ideal variation for testing how firm flexibility shapes responses to pure uncertainty. This represents a substantial shock to the pricing environment, even if not visually dramatic compared to COVID’s extreme spike.

The consistency of the results across these approaches—each addressing different endogeneity concerns—strengthens the causal interpretation of the ‘curse of flexibility’.

2.1.3. Summary Statistics

Figure 1 displays the evolution of price uncertainty from 2003-2021. The summary statistics in Table 1 show that my standardized volatility measure has mean zero and unit standard deviation by construction, with substantial variation across the sample period.

While visual inspection of Figure 1 may suggest modest variation except during COVID and, therefore, modest economic implications, this perception is misleading. This pattern—moderate aggregate variation generating large heterogeneous responses—is precisely what the curse of flexibility predicts.

My identification strategy relies on cross-sectional heterogeneity in price-setting flexibility. The substantial dispersion in adjustment frequencies (mean = 0.296, SD = 0.197) reveals

³This aligns with forward-looking models where firms care about future variance. In the third-order approximation of the Phillips curve, inflation depends on next period’s variance – precisely the behavior my persistence structure supports. The multi-period predictability provides the rational foundation for firms to value the option of waiting during high-dispersion episodes.

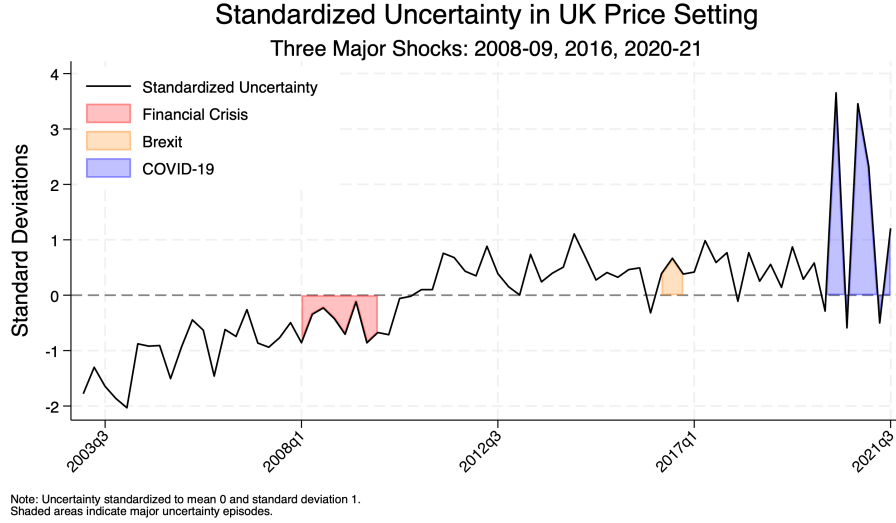


Figure 1: Standardized Uncertainty in UK Price Setting, 2003-2021

Notes: This figure plots the standardised weighted variance of quarterly price changes (mean = 0, standard deviation = 1) from 2003 to 2021. Shaded areas indicate three major uncertainty episodes: the Financial Crisis (2008-09), Brexit referendum (2016), and COVID-19 pandemic (2020-21).

Table 1: Summary Statistics

	Mean	Std. Dev.	Median
Price Change (%)	0.44	10.92	0.00
Adjustment Frequency	0.296	0.197	0.267
Lagged Volatility (std.)	-0.008	0.997	0.102
Lagged Skewness (std.)	-0.002	1.000	-0.022

Notes: Summary statistics for the regression sample from 2003Q2-2021Q3 ($N = 7,427,514$ observations). Price Change is the quarterly log price change in percent, winsorised at the 1st and 99th percentiles. Adjustment Frequency is calculated for each product-store combination as the fraction of quarters in which the unit changes its regular price: $\text{freq}_i = (\text{number of price changes})_i / (\text{total observations})_i$. This unit-level measure ranges from 0 (never adjusts) to 1 (adjusts every quarter). Lagged Volatility and Skewness are standardised to have mean 0 and standard deviation 1 and are measures of the cross-sectional price change distribution.

pronounced differences: units at the 80th percentile adjust 46% of the time, compared to just 12.5% for those at the 20th percentile—a 3.7-fold difference. This heterogeneity provides the requisite variation to identify differential responses to uncertainty shocks across the flexibility spectrum.

3. Empirical Results

Firms facing uncertainty must decide both whether to adjust prices (extensive margin) and by how much (intensive margin). The curse of flexibility primarily operates through the extensive margin—flexible firms reduce their adjustment frequency under uncertainty. Table 2 first examines the overall price adjustment response, combining both margins. I then decompose these effects into extensive and intensive margins.

Table 2 presents the main results examining how price adjustment responds to uncertainty and skewness. The baseline specification in column (1) estimates:

$$\Delta p_{ijt} = \alpha_i + \beta_1 \tilde{\sigma}_{t-1} + \beta_3 \Delta p_{ijt-1} + \gamma X_{ijt} + \epsilon_{ijt} \quad (8)$$

Column (2) extends this to test for heterogeneous responses by firm flexibility.

$$\begin{aligned} \Delta p_{ijt} = & \alpha_i + \beta_1 \tilde{\sigma}_{t-1} + \beta_2 \widetilde{\text{Skew}}_{t-1} + \beta_3 \tilde{\sigma}_{t-1} \times \text{freq}_i \\ & + \beta_4 \widetilde{\text{Skew}}_{t-1} \times \text{freq}_i + \beta_5 \Delta p_{ijt-1} + \gamma X_{ijt} + \epsilon_{ijt} \end{aligned} \quad (9)$$

where Δp_{ijt} is the price change excluding sales, α_i are firm fixed effects, and X_{ijt} includes sectoral and aggregate inflation.

The results reveal striking heterogeneity in how firms respond to uncertainty. Column (1) presents the baseline specification with average effects across all firms (Equation 8). Uncertainty has a negative and statistically significant effect on price adjustment (-0.032), suggesting that the average firm reduces the magnitude of its price changes by 3.2 basis points when uncertainty increases by one standard deviation.

However, this average effect masks substantial heterogeneity. Column (2) introduces an interaction between uncertainty and firm flexibility (adjustment frequency) to test whether firms with different menu costs respond differently to uncertainty. The results are dramatic: the direct effect of uncertainty becomes positive and large (0.094), while the interaction term

Table 2: Effect of Uncertainty on Price Adjustment

	(1) Baseline	(2) Uncertainty Interaction	(3) Full Model
Uncertainty ($\tilde{\sigma}_{t-1}$)	-0.0324*** (0.0054)	0.0942*** (0.0093)	0.0882*** (0.0093)
Skewness ($\widetilde{\text{Skew}}_{t-1}$)			-0.0389*** (0.0068)
Uncertainty \times freq		-0.4208*** (0.0360)	-0.3932*** (0.0359)
Skewness \times freq			0.3939*** (0.0296)
Lagged Price Change	-0.2594*** (0.0007)	-0.2595*** (0.0007)	-0.2597*** (0.0007)
Sectoral Inflation	0.9718*** (0.0057)	0.9718*** (0.0057)	0.9708*** (0.0057)
Aggregate Inflation	0.2379*** (0.0126)	0.2374*** (0.0126)	0.2301*** (0.0126)
Constant	0.0881*** (0.0054)	0.0889*** (0.0055)	0.0927*** (0.0054)
Firm FE	Yes	Yes	Yes
Observations	6,632,240	6,632,240	6,632,240
R^2 (within)	0.077	0.077	0.077
Number of Firms	432,915	432,915	432,915

Notes: The dependent variable is the percentage price change excluding sales. Sample period: 2003Q2-2021Q3. All specifications include firm fixed effects. Uncertainty and skewness are standardised lagged weighted moments of the price change distribution. freq is the firm's historical frequency of price adjustment. Standard errors clustered at the firm level in parentheses.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

is strongly negative (-0.421). This sign reversal indicates that the average negative effect in column (1) is actually a composition of opposing responses across firm types.

To illustrate the economic magnitudes, consider two firms at different points in the flexibility distribution. For a sticky-price firm ($\text{freq} = 0.2$), the total effect of uncertainty is $0.094 + (-0.421 \times 0.2) = 0.010$, which is positive. In contrast, for a flexible-price firm ($\text{freq} = 0.8$), the total effect is $0.094 + (-0.421 \times 0.8) = -0.243$. This means flexible firms reduce their price adjustments by 24.3 basis points per standard deviation increase in uncertainty, while sticky firms show an increase. The differential response of 25.3 basis points between flexible and sticky firms is both statistically significant and economically substantial.

Since uncertainty is measured as the second moment of the price change distribution, it may partially capture higher-order distributional effects. Column (3) therefore includes skewness (the third moment) and its interaction with flexibility as a robustness check. The inclusion of skewness controls leaves the uncertainty results virtually unchanged: the direct effect remains positive (0.088) and the interaction remains strongly negative (-0.393). This stability confirms that the heterogeneous uncertainty effects are not driven by omitted higher moments.

The skewness results themselves reveal an opposite pattern. The direct effect is negative (-0.039), but the interaction with flexibility is positive (0.394). For the same sticky firm ($\text{freq} = 0.2$), the total skewness effect is $-0.039 + (0.394 \times 0.2) = 0.040$, while for the flexible firm ($\text{freq} = 0.8$), it is $-0.039 + (0.394 \times 0.8) = 0.276$. This indicates that flexible firms increase their price adjustments by 27.6 basis points when the price distribution becomes more right-skewed, while sticky firms barely respond.

The contrasting signs of the uncertainty and skewness interactions reveal distinct adjustment patterns. Flexible firms appear to follow a cautious strategy during periods of high volatility (reducing adjustment magnitudes) but respond aggressively to directional price pressures captured by positive skewness. This asymmetric response suggests that flexible and sticky firms differ not only in their adjustment frequency but also in how they process different types of aggregate shocks.

3.1. Extensive Margin

I now examine the extensive margin—whether firms adjust their prices at all. I estimate a linear probability model:

$$\begin{aligned} \mathbb{I}(\Delta p_{ijt} \neq 0) = & \alpha_i + \beta_1 \tilde{\sigma}_{t-1} + \beta_2 \widetilde{\text{Skew}}_{t-1} + \beta_3 \tilde{\sigma}_{t-1} \times \text{freq}_i \\ & + \beta_4 \widetilde{\text{Skew}}_{t-1} \times \text{freq}_i + \beta_5 \Delta p_{ijt-1} + \gamma X_{ijt} + \epsilon_{ijt} \end{aligned} \quad (10)$$

where $\mathbb{I}(\Delta p_{ijt} \neq 0)$ is an indicator for any price adjustment. All specifications reported in this section include firm fixed effects and control for lagged price changes, sectoral inflation, and aggregate inflation.

Table 3: Decomposing the Curse of Flexibility: Intensive vs Extensive Margins

	(1)	(2)
	Prob(Adjust)	Size Adjust
Uncertainty ($\tilde{\sigma}_{t-1}$)	−0.009*** (0.000)	0.463*** (0.030)
Uncertainty × freq	−0.018*** (0.001)	−0.325*** (0.066)
Controls	Yes	Yes
Firm FE	Yes	Yes
Observations	6,851,936	1,973,123
R^2	0.007	0.005

Notes: Column (1) shows the extensive margin using a linear probability model for any price adjustment. Column (2) shows the intensive margin using absolute price changes conditional on adjustment. All specifications include controls for skewness and its interaction with flexibility, firm fixed effects, lagged price changes, and inflation controls. Sample period: 2003Q2-2021Q3. Standard errors clustered at the firm level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Column (1) of Table 3 demonstrates the curse of flexibility. Uncertainty has a negative and statistically significant effect on price adjustment probability (−0.009). Importantly, the interaction coefficient of −0.018 indicates that flexible firms reduce their adjustment probability more than sticky firms when uncertainty increases by one standard deviation. For a sticky firm ($\text{freq} = 0.2$), the total effect is $-0.009 + (-0.018 \times 0.2) = -0.0126$, or a 1.26 percentage point reduction. For a flexible firm ($\text{freq} = 0.8$), the effect is $-0.009 + (-0.018 \times 0.8) = -0.0234$, or a 2.34 percentage point reduction. Flexible firms thus reduce their adjustment probability nearly twice as much as sticky firms.

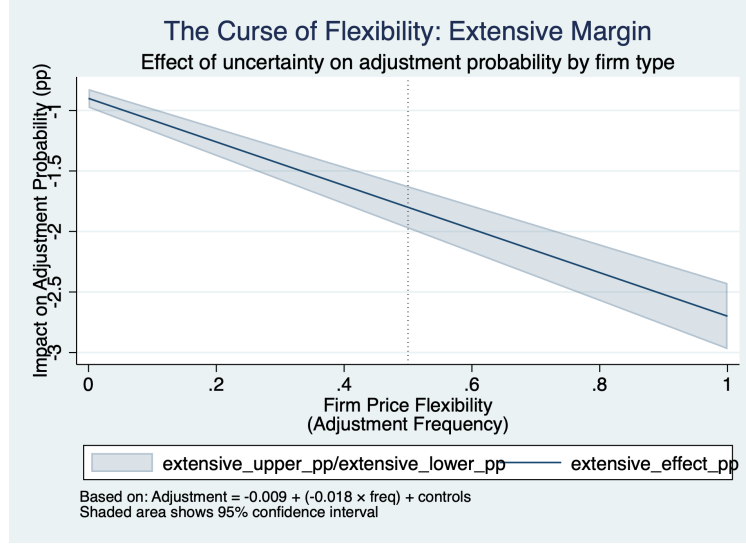


Figure 2: The Curse of Flexibility

Notes: This figure shows how the effect of uncertainty on price adjustment varies with firm flexibility. The solid line represents the marginal effect of a one standard deviation increase in uncertainty on the probability of price adjustment, calculated as $-0.009 - 0.018 \times \text{flexibility}$. Firm flexibility is measured as the historical frequency of price adjustment. The shaded area shows 95% confidence intervals. All firms experience reduced adjustment under uncertainty, but the effect is strongest for flexible firms.

These results clearly show the curse of flexibility: the most flexible firms cut their adjustment probability the most when uncertainty rises.

Figure 2 provides visual evidence of the curse of flexibility. The downward-sloping relationship demonstrates that firms with greater price flexibility reduce their adjustment probability most when uncertainty rises. The curse intensifies continuously with flexibility, reaching its maximum effect among the most flexible firms. This heterogeneous response is economically meaningful: flexible firms (frequency=0.8) reduce adjustment by 1.08 percentage points more than sticky firms (frequency=0.2).

3.2. Identification with pre-determined measures

A key concern is that flexibility measures might be endogenous to uncertainty episodes: it might be that uncertainty itself causes firms to change their adjustment behavior, creating a spurious correlation. Specifically, the *curse* might appear simply because firms adjusted frequently during a particular uncertainty shock, then subsequently froze. This would generate the observed negative interaction without any true curse of flexibility.

To address this concern, I measure flexibility using two pre-determined approaches:

Pre-2005 Flexibility: I calculate each firm’s adjustment frequency using only data from 2005 and earlier, before the major uncertainty episodes in the sample (Financial Crisis, Brexit, COVID-19). This creates a clean measure of inherent flexibility uncontaminated by responses to large shocks.

Calm-Period Flexibility: I measure flexibility during the lowest uncertainty tercile. This captures firm types under *normal* conditions.

Both approaches mimic a quasi-experimental design: firms are assigned flexibility types before treatment (uncertainty shocks), allowing causal interpretation of differential responses.

Table 4: Robustness: Alternative Flexibility Measures

	(1)	(2)	(3)
	Main	Pre-2005	Calm Period
	Specification	Flexibility	Flexibility
Uncertainty ($\tilde{\sigma}_{t-1}$)	−0.009*** (0.000)	0.036*** (0.001)	−0.015*** (0.001)
Uncertainty \times freq	−0.018*** (0.001)	−0.135*** (0.002)	−0.061*** (0.002)
Controls	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes
Observations	6,851,936	3,490,393	2,303,301
R^2	0.007	0.013	0.016

Notes: Column 1 reproduces the main specification. Column 2 uses flexibility measured from pre-2005 data. Column 3 uses flexibility measured during calm periods (lowest uncertainty tercile) and estimates effects during volatile periods (upper two terciles). All specifications include controls for skewness and its interaction with flexibility, firm fixed effects, lagged price changes, and inflation controls. Standard errors clustered at the firm level. *** $p < 0.01$.

Table 4 confirms the curse’s robustness. Both alternative flexibility measures yield even stronger curse effects. Using pre-2005 flexibility (column 2), the interaction coefficient of -0.135 is 7.5 times larger in magnitude than in the main specification. Using calm-period flexibility (column 3), the coefficient of -0.061 is 3.4 times larger. This pattern—stronger effects with cleaner measurement—suggests the main results understate the true curse due to endogenous flexibility classification.

3.3. Intensive Margin

For the intensive margin, I estimate:

$$\begin{aligned}
|\Delta p_{ijt}| = & \alpha_i + \beta_1 \tilde{\sigma}_{t-1} + \beta_2 \widetilde{\text{Skew}}_{t-1} + \beta_3 \tilde{\sigma}_{t-1} \times \text{freq}_i \\
& + \beta_4 \widetilde{\text{Skew}}_{t-1} \times \text{freq}_i + \beta_5 \Delta p_{ijt-1} + \gamma X_{ijt} + \epsilon_{ijt}
\end{aligned} \tag{11}$$

conditional on $\Delta p_{ijt} \neq 0$. The **intensive margin** (columns 2 of Table 3) shows a complementary pattern. While flexible firms adjust less frequently under uncertainty, when they do adjust, they also make smaller price changes. The negative interaction in column 2 (-0.325) means flexible firms reduce the absolute size of their adjustments by 0.325 percentage points more than sticky firms. Column 2 confirms this pattern specifically for price increases: flexible firms reduce the size of their price increases by 0.27 percentage points more than sticky firms (interaction of -0.267). This suggests a general pattern of caution rather than the standard model prediction of “wait longer, adjust more.”

Together, these results along with the ones for extensive margin paint a picture of flexible firms becoming cautious across multiple dimensions when uncertainty rises: they adjust less frequently, make smaller changes when they do adjust, and particularly avoid price increases. This multi-dimensional caution contradicts standard menu cost models, which predict that higher uncertainty should make flexibility more valuable.

4. The Curse of Flexibility During Brexit: a Natural Experiment

The June 2016 Brexit referendum provides an ideal quasi-experimental setting to test how firm heterogeneity shapes responses to uncertainty. Unlike episodes involving operational disruptions—natural disasters, or financial crises with bank failures—UK businesses continued normal operations while facing elevated uncertainty about future trading relationships. This isolates voluntary economic responses from mechanical constraints, providing a cleaner test of the curse of flexibility.

Table 5 examines the post-referendum period. The pooled specification (column 1) shows an interaction coefficient of -0.0295 ($p < 0.01$), indicating that each unit increase in flexibility is associated with an additional 2.95 percentage point reduction in adjustment probability when uncertainty rises by one standard deviation. For firms with frequency=0.8 versus frequency=0.2, this implies a differential response of 1.8 percentage points. Column 2 focuses on Q2 2017-2019, when uncertainty and skewness are nearly orthogonal ($\rho = 0.064$), providing

particularly clean identification. Here the curse strengthens, with an interaction coefficient of -0.0336, implying a 2.0 percentage point differential between firms with frequency=0.8 versus frequency=0.2.

Table 5: The Curse of Flexibility During Brexit

	(1) Pooled Post-Brexit	(2) Q2 2017-2019
Uncertainty ($\tilde{\sigma}_{t-1}$)	-0.0140*** (0.0018)	-0.0158*** (0.0018)
Skewness ($\widetilde{\text{Skew}}_{t-1}$)	0.0096*** (0.0006)	0.0070*** (0.0009)
Uncertainty \times freq	-0.0295*** (0.0059)	-0.0336*** (0.0061)
Skewness \times freq	0.0212*** (0.0021)	0.0408*** (0.0031)
Corr(σ , Skew)	-0.283	0.064
Controls	Yes	Yes
Firm FE	Yes	Yes
Observations	1,244,844	941,199
R^2 (within)	0.004	0.004
Number of Firms	180,232	163,838

Notes: The dependent variable is an indicator for price adjustment. Column 1 pools all post-Brexit quarters (Q3 2016-2019). Column 2 focuses on Q2 2017-2019 when identification is cleanest. All specifications include firm fixed effects, lagged price changes, and sectoral inflation controls. Aggregate inflation omitted due to collinearity in quarterly subsamples. Standard errors clustered at the firm level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Flexible firms (frequency=0.8) reduce adjustment by 3.8 percentage points while sticky firms (frequency=0.2) reduce by 2.0 percentage points—a differential of 1.8 percentage points. While this differential may appear modest, it represents a complete reversal of theoretical predictions. Menu cost models unambiguously predict that lower adjustment costs should amplify responses to shocks. Finding the opposite—that flexible firms become more cautious—challenges the core logic of price-setting models, regardless of magnitude.

The Brexit evidence yields an important methodological insight: the correlation structure between uncertainty and higher moments can mask heterogeneous firm responses. During

Q3 2016-Q1 2017, uncertainty and skewness were highly correlated ($\rho = -0.806$), making their effects statistically indistinguishable. However, in Q2 2017-2019, when these moments became nearly orthogonal ($\rho = 0.064$), the curse of flexibility emerges cleanly and strengthens from -0.0295 to -0.0336. This pattern may help reconcile different findings in the uncertainty literature—studies examining operational disruptions capture combined effects of multiple shocks, while our approach isolates the specific channel through which pure uncertainty affects heterogeneous firms when higher moments can be separated. The curse’s emergence specifically when identification is cleanest—and its consistency across multiple identification strategies—confirms it reflects firm optimisation rather than statistical confounding.

5. Comparing Uncertainty and Monetary Policy Responses

The economic significance of the curse of flexibility becomes clear when contrasted with a well-understood shock: monetary policy. If flexibility is truly a ‘curse’ under uncertainty but an ‘advantage’ under standard shocks, this would demonstrate a fundamental context-dependency in firm behavior that existing models miss. For monetary policy, I employ high-frequency identification using the UK Monetary Policy Event-Study Database developed by Braun et al. (2025). This dataset extracts monetary surprises from asset price movements in narrow windows around Bank of England announcements, isolating exogenous variation in the monetary policy stance from endogenous responses to economic conditions. The database covers Monetary Policy Committee decisions and Inflation Report releases from 2000-2021, providing clean variation in monetary stance. I standardize monetary policy shocks to have unit variance for comparability with uncertainty shocks.

The results reveal starkly opposite patterns. Under uncertainty (columns 1-2), the interaction coefficients are consistently negative and significant (-0.018 and -0.0295), implying that flexible firms (freq=0.8) reduce their adjustment probability 1.08 and 1.8 percentage points more than sticky firms (freq=0.2) in the full sample and Brexit period respectively. In contrast, monetary policy shocks (column 3) generate the standard theoretical prediction: a positive interaction coefficient (0.0152) means flexible firms increase adjustment by 0.91 percentage points more than sticky firms.

This reversal is both statistically robust and economically meaningful. During the Brexit

Table 6: The Curse vs. Advantage of Flexibility: Uncertainty and Monetary Policy

	(1) Full Sample Uncertainty	(2) Brexit Uncertainty	(3) Monetary Policy Shock
Shock Variable	−0.009*** (0.000)	−0.0140*** (0.0018)	−0.0032*** (0.0003)
Shock \times freq	−0.018*** (0.001)	−0.0295*** (0.0059)	0.0152*** (0.0009)
Economic Effects:			
Flexible firms (freq=0.8)	−2.34pp	−3.76pp	+0.90pp
Sticky firms (freq=0.2)	−1.26pp	−1.99pp	−0.02pp
Differential	−1.08pp	−1.77pp	+0.92pp
Interpretation	Curse of Flexibility	Curse of Flexibility	Advantage of Flexibility
Controls	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes
Observations	6,851,936	1,244,844	7,007,661

Notes: Dependent variable is probability of price adjustment. Column (1) uses full sample 2003-2021, Column (2) focuses on Brexit period (Q3 2016-2019), Column (3) examines monetary policy shocks. All specifications include controls for skewness interactions, firm fixed effects, lagged price changes, and inflation controls. Standard errors clustered at firm level. *** $p < 0.01$.

period, uncertainty generates differential responses of 1.8 percentage points (pooled) to 2.0 percentage points (Q2 2017-2019) between flexible and sticky firms, while monetary policy generates a 0.91 percentage point differential—opposite in sign. This demonstrates that firm heterogeneity is not a fixed characteristic but depends fundamentally on the nature of the shock.

The implications for policy transmission are profound. Monetary policy relies on flexible firms’ responsiveness, yet uncertainty weakens precisely this channel. When both shocks coincide, uncertainty does not just moderate policy transmission—it can weaken or even reverse monetary policy’s intended effects, as flexible firms that normally amplify policy transmission instead become the most cautious. This state-dependence suggests that models assuming constant firm types miss crucial variation in how different shocks propagate through heterogeneous firms.

6. Heterogeneous Menu Costs Model

Having established the curse of flexibility empirically, I now develop a theoretical framework to explain this paradox. While standard menu cost models predict that flexibility always enhances responsiveness to shocks, the data show the opposite under uncertainty. To reconcile theory with evidence, I build a model with heterogeneous menu costs and demonstrate how this heterogeneity creates differential exposure to uncertainty between adjusters and non-adjusters—a mechanism that has been overlooked in the literature.

The model is deliberately structured to isolate the differential exposure mechanism. While I present a full dynamic model with persistent shocks, the key analytical results are derived by analysing firm behavior conditioning on realised uncertainty levels, treating within-period shocks as i.i.d. This approach—examining behavior within uncertainty regimes rather than transitions between them—allows analytical characterisation of why flexible firms reduce adjustment under high uncertainty.

The key innovation is recognising that menu cost heterogeneity doesn’t just determine adjustment frequencies; it fundamentally alters how firms experience uncertainty. Firms with different menu costs face different trade-offs between the risks of adjusting (symmetric exposure from reset prices) versus not adjusting (asymmetric exposure from existing posi-

tions). This differential exposure mechanism generates non-monotonic responses that can reverse the standard predictions when uncertainty is sufficiently high.

In the model there is a continuum of firms and firms differ in their price adjustment costs. Each firm i draws a menu cost parameter ψ_i from a continuous distribution $F(\psi)$ with support $[\underline{\psi}, \bar{\psi}]$, where $0 < \underline{\psi} < \bar{\psi} < \infty$. The menu cost represents the labor units required to implement a price change, implying a real adjustment cost of $\psi_i \cdot w_t$, where $w_t = W_t/P_t$ denotes the real wage. These menu costs are permanent firm characteristics that do not vary over time.

The economy is subject to aggregate TFP shocks. Log productivity follows an AR(1) process:

$$\log A_t = \rho_A \log A_{t-1} + \sigma_t \varepsilon_t \quad (12)$$

where $\varepsilon_t \sim N(0, 1)$ and, crucially, the uncertainty of TFP shocks varies over time according to

$$\sigma_t^2 = \bar{\sigma}^2 + \rho_\sigma(\sigma_{t-1}^2 - \bar{\sigma}^2) + \eta_t \quad (13)$$

where $\bar{\sigma}^2$ is the unconditional variance, $\rho_\sigma \in [0, 1)$ governs persistence, and η_t represents uncertainty shocks. The timing within each period t unfolds as follows. First, firms observe the aggregate state $\Omega_t = \{A_t, P_t, Y_t, \sigma_t^2\}$, which includes the current level of TFP uncertainty. Second, each firm decides whether to adjust its price. Third, firms that choose to adjust pay their menu cost $\psi_i \cdot w_t$ and set a new price. Fourth, production and consumption take place given the realised prices. Finally, TFP shocks realize, determining next period's productivity A_{t+1} .

In the rest of this section, I outline the building blocks of the model. I start by describing the household sector, which is standard.

6.1. Households

The economy is populated by a representative household with preferences over consumption streams given by

$$U = \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t u(C_t) \quad (14)$$

where $\beta \in (0, 1)$ is the discount factor and C_t represents a CES aggregate of differentiated goods:

$$C_t = \left[\int_0^1 c_{it}^{\frac{\epsilon-1}{\epsilon}} di \right]^{\frac{\epsilon}{\epsilon-1}} \quad (15)$$

with $\epsilon > 1$ denoting the elasticity of substitution across varieties. Standard optimisation yields the demand function for each variety i :

$$c_{it} = \left(\frac{P_{it}}{P_t} \right)^{-\epsilon} C_t \quad (16)$$

where $P_t = \left[\int_0^1 P_{it}^{1-\epsilon} di \right]^{\frac{1}{1-\epsilon}}$ is the aggregate price index.

6.2. Firms

Production is carried out by a continuum of monopolistically competitive firms indexed by $i \in [0, 1]$. Each firm operates a linear production technology $y_{it} = A_t n_{it}$, where A_t represents aggregate productivity and n_{it} is labor input. Given the household's demand function, firm i faces the demand curve

$$y_{it} = \left(\frac{P_{it}}{P_t} \right)^{-\epsilon} Y_t \quad (17)$$

where Y_t denotes aggregate output. The nominal marginal cost common to all firms is $MC_t = W_t/A_t$, where W_t is the nominal wage. Firms differ only in their menu cost parameters ψ_i . In equilibrium, firms with identical menu costs behave identically, allowing us to index firms by their menu cost type. This heterogeneity generates endogenous price stickiness that varies across the menu cost distribution—firms with high ψ adjust infrequently (sticky prices), while firms with low ψ adjust more often (flexible prices). As we will show, this seemingly natural ordering is reversed when uncertainty is high, leading to the curse of flexibility.

6.3. The Firm's Problem

Each firm i chooses its price to maximise the expected discounted value of profits. The firm's state is characterised by its current price P_{it} (or equivalently, the log relative price $\hat{p}_{it} \equiv \log(P_{it}/P_t)$) and the aggregate state Ω_t . The Bellman equation is:

$$V(\hat{p}_{it}, \Omega_t; \psi_i) = \max \{V^n(\hat{p}_{it}, \Omega_t; \psi_i), \quad V^a(\Omega_t; \psi_i) - \psi_i w_t\} \quad (18)$$

where V^n denotes the value of not adjusting and V^a denotes the value after adjustment.

If the firm does not adjust, it earns current profits and its relative price erodes with inflation:

$$V^n(\hat{p}_{it}, \Omega_t; \psi_i) = \Pi(\hat{p}_{it}, \Omega_t) + \beta \mathbb{E}_t[V(\hat{p}_{it} - \pi_{t+1}, \Omega_{t+1}; \psi_i)] \quad (19)$$

where $\pi_{t+1} \equiv \log(P_{t+1}/P_t)$ is the inflation rate and the expectation is taken over the joint distribution of $(A_{t+1}, \sigma_{t+1}^2)$.

If the firm adjusts, it chooses the optimal reset price \hat{p}_t^R :

$$V^a(\Omega_t; \psi_i) = \max_{\hat{p}} \{\Pi(\hat{p}, \Omega_t) + \beta \mathbb{E}_t[V(\hat{p} - \pi_{t+1}, \Omega_{t+1}; \psi_i)]\} \quad (20)$$

The profit function in terms of log relative prices is:

$$\Pi(\hat{p}_{it}, \Omega_t) = Y_t [\exp((1 - \epsilon)\hat{p}_{it}) - mc_t \exp(-\epsilon\hat{p}_{it})] \quad (21)$$

where $mc_t = W_t/(P_t A_t)$ is the real marginal cost.

6.4. Optimal Reset Price

When a firm pays the menu cost to adjust, it sets a forward-looking price that accounts for future inflation and uncertainty. The first-order condition for the optimal reset price \hat{p}_t^R balances current marginal profit against discounted future marginal value:

$$\Pi_{\hat{p}}(\hat{p}_t^R, \Omega_t) + \beta \mathbb{E}_t[V_{\hat{p}}(\hat{p}_t^R - \pi_{t+1}, \Omega_{t+1})] = 0 \quad (22)$$

Proposition 1 (Optimal Reset Price). *The optimal reset price up to second order is:*

$$\hat{p}_t^R = \hat{p}_t^* + \alpha_1 \mathbb{E}_t[\pi_{t+1}] + \alpha_2 \text{Var}_t(\pi_{t+1}) \quad (23)$$

where \hat{p}_t^* is the frictionless optimal price and the coefficients are:

$$\alpha_1 = \frac{\beta(1 + \rho)}{1 + \beta} \quad (24)$$

$$\alpha_2 = \frac{\beta(\epsilon - 2)(1 + \rho)^2}{2(1 + \beta)} \quad (25)$$

with $\rho \in [0, 1]$ measuring the pass-through of inflation to marginal costs and $\epsilon > 1$ the demand elasticity.

Proof: See Appendix B.

This pricing rule reveals how forward-looking firms respond to an uncertain environment. There are two channels: (1) Expected Inflation Adjustment (α_1 term) and (2) Uncertainty Premium (α_2 term). Firms raise prices above the current optimum to compensate for expected erosion from inflation. With full pass-through ($\rho = 1$) and patient firms ($\beta \rightarrow 1$), they offset approximately half of expected inflation. This partial adjustment reflects the trade-off between current profit losses from high prices and future losses from eroded prices. The second channel suggests that when demand is convex in prices ($\epsilon > 2$), inflation variance raises reset prices further. This precautionary behavior emerges because profit losses from prices being too low exceed gains from prices being too high. The variance effect increases with the degree of pass-through and demand convexity.

It is interesting to note that when $\epsilon = 2$, the variance term vanishes and we recover the certainty-equivalent pricing of Golosov and Lucas (2007). Empirical estimates typically find $\epsilon \in [4, 11]$, well above this threshold, confirming the variance premium matters in practice. With no pass-through ($\rho = 0$), only demand-side inflation matters for pricing. As $\beta \rightarrow 0$ (myopic firms), both effects disappear and firms set static optimal prices.

6.5. The Adjustment Decision

Having characterised the optimal reset price, I now analyse the adjustment decision. A firm will reset its price when the net gain exceeds the menu cost:

$$G(x_{it}, \Omega_t; \psi_i) \equiv \underbrace{V(\hat{p}_t^R, \Omega_t) - \psi_i w_t}_{\text{Value if adjust}} - \underbrace{V(\hat{p}_{it}, \Omega_t)}_{\text{Value if don't adjust}} > 0 \quad (26)$$

where $x_{it} \equiv \hat{p}_{it} - \hat{p}_t^*$ is the price gap. This generates an (S, s) policy with time-varying thresholds $[\underline{x}(\Omega_t; \psi_i), \bar{x}(\Omega_t; \psi_i)]$.

It is helpful to begin the discussion by considering the standard menu cost model. In conventional menu cost models, higher uncertainty σ_t^2 increases adjustment frequency for all firms:

$$\frac{\partial \Pr(\text{adjust}|\psi_i)}{\partial \sigma^2} > 0 \quad \forall \psi_i \quad (27)$$

This occurs because uncertainty widens the distribution of price gaps, pushing more firms beyond their adjustment thresholds. Under higher uncertainty shocks to costs and demand can be larger. As a consequence, price gaps x_{it} disperse more quickly and more firms hit their adjustment thresholds per period. Adjustment frequency rises uniformly across all menu cost types

This prediction is intuitive and has shaped how economists think about price adjustment in volatile times. It suggests that flexible-price firms (low ψ_i) should be especially responsive to uncertainty since they face lower barriers to adjustment.

However, as I discussed earlier, the data strongly reject this prediction for flexible firms. This failure of standard theory is not a minor deviation—it is a complete reversal of the predicted relationship. Understanding this puzzle requires rethinking how adjustment costs interact with uncertainty.

7. The Mechanism Behind the Curse of Flexibility

The standard model predicts monotonic responses: lower menu costs always lead to more adjustment under uncertainty. I now show how differential exposure between adjusters and non-adjusters can reverse this relationship. The key insight is that adjustment thresholds respond endogenously to uncertainty in a non-monotonic way that depends on menu costs.

7.1. The Firm's Dynamic Problem

For analytical tractability, I now analyse firm behavior conditioning on realised uncertainty levels, treating inflation shocks as i.i.d. within each uncertainty regime. This approach allows me to derive closed-form solutions that clearly demonstrate the differential exposure mechanism while maintaining the key economic insights from the full model.

Consider a firm with menu cost ψ facing a price gap $x = \hat{p} - \hat{p}^*$ from its static optimum. The firm's state is characterised by the triple (x, σ^2, ψ) , where σ^2 represents the current level of aggregate uncertainty. Each period, the firm chooses between adjusting its price or maintaining the status quo.

The Bellman equation for this problem is:

$$V(x, \sigma^2; \psi) = \max\{V^n(x, \sigma^2; \psi), V^a(x, \sigma^2; \psi)\} \quad (28)$$

where V^n denotes the value of not adjusting and V^a denotes the value of adjusting.

The value of not adjusting is given by:

$$V^n(x, \sigma^2; \psi) = \pi(x) + \beta \mathbb{E}[V(x - \pi', \sigma^{2'}; \psi) | \sigma^2] \quad (29)$$

where $\pi(x)$ represents current-period profits as a function of the price gap, π' is next period's inflation rate, and the expectation is taken over the joint distribution of $(\pi', \sigma^{2'})$ conditional on current uncertainty σ^2 .

The value of adjusting is:

$$V^a(x, \sigma^2; \psi) = \pi(0) - \psi w + \beta \mathbb{E}[V(-\pi', \sigma^{2'}; \psi) | \sigma^2] \quad (30)$$

where $\pi(0)$ represents profits at the optimal price and ψw is the real menu cost.

The key difference between the two cases is that the non-adjuster carries forward her current gap x , which is then augmented by inflation, while the adjuster resets to zero but immediately faces erosion from future inflation.

7.2. The Adjustment Decision

Define the gain from adjustment as:

$$G(x, \sigma^2; \psi) \equiv V^a(x, \sigma^2; \psi) - V^n(x, \sigma^2; \psi) \quad (31)$$

The firm adjusts when $G(x, \sigma^2; \psi) > 0$. The adjustment threshold $\bar{x}(\sigma^2, \psi)$ is implicitly defined by:

$$G(\bar{x}(\sigma^2, \psi), \sigma^2; \psi) = 0 \quad (32)$$

The adjustment threshold \bar{x} now depends on both the current uncertainty level σ^2 and the menu cost ψ . The key theoretical result will show that $\partial \bar{x} / \partial \sigma^2 > 0$ for low ψ (flexible firms) but $\partial \bar{x} / \partial \sigma^2 \approx 0$ for high ψ (sticky firms).

7.3. Effect of Uncertainty on the Adjustment Threshold

I now analyse how changes in uncertainty affect the firm's adjustment behavior. The key object of interest is $\partial \bar{x} / \partial \sigma^2$ —how does the adjustment threshold respond to uncertainty? I focus on positive price gaps $x > 0$; by symmetry of the profit function, the analysis applies

equally to negative gaps with threshold $-\bar{x}$. This derivative will reveal whether firms become more or less willing to tolerate price deviations when facing greater uncertainty.

Since the threshold $\bar{x}(\sigma^2, \psi)$ is implicitly defined by the indifference condition, we have:

$$G(\bar{x}(\sigma^2, \psi), \sigma^2; \psi) = 0 \quad (33)$$

I can apply the implicit function theorem. The threshold must continuously satisfy this indifference condition, so any change in σ^2 must be accompanied by a change in \bar{x} that maintains $G = 0$.

Totally differentiating the indifference condition with respect to σ^2 :

$$\frac{d}{d\sigma^2}[G(\bar{x}(\sigma^2, \psi), \sigma^2; \psi)] = 0 \quad (34)$$

Since G depends on σ^2 both directly and through $\bar{x}(\sigma^2)$, I apply the chain rule:

$$\left. \frac{\partial G}{\partial x} \right|_{x=\bar{x}} \cdot \frac{d\bar{x}}{d\sigma^2} + \left. \frac{\partial G}{\partial \sigma^2} \right|_{x=\bar{x}} = 0 \quad (35)$$

Note that I use $d\bar{x}/d\sigma^2$ (total derivative) because \bar{x} is a function of σ^2 , while I use partial derivatives for G because I'm holding other arguments fixed.

Solving for the threshold response:

$$\frac{\partial \bar{x}}{\partial \sigma^2} = - \left. \frac{\partial G / \partial \sigma^2}{\partial G / \partial x} \right|_{x=\bar{x}} \quad (36)$$

This expression reveals that the threshold's response to uncertainty depends on two key components:

1. **Numerator** ($\partial G / \partial \sigma^2$): How does uncertainty directly affect the gain from adjustment? If uncertainty makes adjustment less attractive ($\partial G / \partial \sigma^2 < 0$), this contributes to a higher threshold.
2. **Denominator** ($\partial G / \partial x$): How does the gain from adjustment change as the price gap increases? At the margin of adjustment, we expect $\partial G / \partial x > 0$ —larger gaps make adjustment more attractive because the profit loss from suboptimal pricing increases.

Since the gain from adjustment increases with the price gap ($\partial G / \partial x > 0$ at the threshold), the threshold response to uncertainty has the opposite sign of $\partial G / \partial \sigma^2$. I will show that

uncertainty reduces the gain from adjustment for flexible firms, causing them to raise their thresholds and adjust less frequently—the curse of flexibility.

7.4. Computing the Effect of Uncertainty on the Gain from Adjustment

This section computes $\partial G/\partial \sigma^2$ and shows why it differs between flexible and sticky firms.

Recall that:

$$G(x, \sigma^2; \psi) = V^a(x, \sigma^2; \psi) - V^n(x, \sigma^2; \psi) \quad (37)$$

Taking the partial derivative with respect to σ^2 :

$$\frac{\partial G}{\partial \sigma^2} = \frac{\partial V^a}{\partial \sigma^2} - \frac{\partial V^n}{\partial \sigma^2} \quad (38)$$

To understand what this means economically: I’m asking how an increase in uncertainty affects the relative attractiveness of adjusting versus not adjusting. If this derivative is negative, it means uncertainty makes adjustment less attractive relative to waiting—the key mechanism behind the curse of flexibility.

7.4.1. Proof of Differential Exposure

Consider a firm with price gap $x > 0$, quadratic profit loss $\Pi(x) = -\frac{\kappa}{2}x^2$, inflation shocks $\pi' \sim N(0, \sigma^2)$, and small adjustment cost ψ .

Key Lemma. For small menu costs ψ , the value function near the adjustment threshold can be approximated as:

$$V(x, \sigma^2) \approx -\frac{\kappa}{2}x^2 - \psi \cdot \mathbb{I}_{|x| > \bar{x}} \quad (39)$$

where $\bar{x} = \sqrt{2\psi/\kappa}$ is the standard menu cost threshold.

This approximation captures the key economic trade-offs facing the firm. The first term, $-\frac{\kappa}{2}x^2$, represents the current-period profit loss from having a suboptimal price gap x . The second term, $\psi \cdot \mathbb{I}_{|x| > \bar{x}}$, represents the expected future adjustment cost, paid only when adjustment is optimal.

Small menu costs imply frequent adjustment, simplifying the value function to current profits and next-period adjustment costs.

Proof of $\partial G/\partial \sigma^2 < 0$. Using the lemma:

$$V^a(\sigma^2) \approx -\frac{\kappa}{2}\sigma^2 - \psi \Pr(|\pi'| > \bar{x}) \quad (40)$$

$$V^n(x, \sigma^2) \approx -\frac{\kappa}{2}(x^2 + \sigma^2) - \psi \Pr(|x - \pi'| > \bar{x}) \quad (41)$$

The key insight is that adjusters and non-adjusters face different exposure to uncertainty. Using the value function approximation, adjusters who reset to $x = 0$ face probability

$$P^a(\sigma) = \Pr(|\pi'| > \bar{x}) = 2\Phi(-\bar{x}/\sigma) \quad (42)$$

of needing to adjust next period. Non-adjusters who maintain gap x face probability

$$P^n(x, \sigma) = \Pr(|x - \pi'| > \bar{x}) = \Phi\left(\frac{x - \bar{x}}{\sigma}\right) + \Phi\left(\frac{-x - \bar{x}}{\sigma}\right) \quad (43)$$

where $\Phi(\cdot)$ is the standard normal CDF.

The crucial step is to establish how sensitive adjusters and non-adjusters are to uncertainty:

Theorem 1 (Sensitivity Comparison and differential exposure). The sensitivity difference $\Delta(\sigma) = \frac{\partial P^a}{\partial \sigma} - \frac{\partial P^n}{\partial \sigma}$ satisfies:

$$\Delta(\sigma) \begin{cases} < 0 & \text{if } \sigma < \sigma^* \text{ (non-adjusters more sensitive)} \\ = 0 & \text{if } \sigma = \sigma^* \text{ (equal sensitivities)} \\ > 0 & \text{if } \sigma > \sigma^* \text{ (adjusters more sensitive)} \end{cases} \quad (44)$$

For $\sigma > \sigma^*$, the difference exhibits an inverted-U shape: it reaches a maximum at some $\sigma_{max} > \sigma^*$ and then decreases as $\Delta(\sigma) \sim 3\bar{x}x^2/\sigma^2$, approaching zero as $\sigma \rightarrow \infty$.

Proof. To establish existence of σ^* , I examine the limiting behavior. For very small σ , non-adjusters are close to the adjustment boundary and become extremely sensitive to uncertainty changes, while adjusters enjoy a buffer zone. Thus $\Delta(\sigma) < 0$ for small σ . For very large σ , adjusters face symmetric exposure to shocks from both directions, while non-adjusters have asymmetric exposure that partially protects them. Thus $\Delta(\sigma) > 0$ for large σ . Since $\Delta(\sigma)$ is continuous and changes from negative to positive, there must exist a unique σ^* where $\Delta(\sigma^*) = 0$. The inverted-U shape for $\sigma > \sigma^*$ arises because the sensitivity advantage of adjusters first increases as we move away from σ^* , but eventually diminishes as extreme uncertainty makes both types equally likely to adjust. See Appendix C for the formal asymptotic analysis. \square

For $\sigma > \sigma^*$, this differential sensitivity creates an differential exposure effect, as the gain

from adjustment decreases with uncertainty:

$$\frac{\partial G}{\partial \sigma^2} = -\frac{\beta\psi}{2\sigma} \left(\frac{\partial P^a}{\partial \sigma} - \frac{\partial P^n}{\partial \sigma} \right) < 0 \quad (45)$$

7.5. Economic Mechanism: From Reset Prices to the Curse

This differential sensitivity result connects directly to the optimal reset price derived in Section 3.4. Recall that when firms adjust, they set prices that include an uncertainty premium:

$$\hat{p}_t^R = \hat{p}_t^* + \alpha_1 \mathbb{E}_t[\pi_{t+1}] + \alpha_2 \text{Var}_t(\pi_{t+1}) \quad (46)$$

The variance term $\alpha_2 \text{Var}_t(\pi_{t+1})$ represents precautionary pricing—firms set higher prices to buffer against future uncertainty. This creates a critical divergence in starting positions:

Adjusters reset to the elevated price \hat{p}_t^R , which includes the variance premium. When uncertainty resolves, they typically find themselves with positive price gaps (prices too high relative to the realized optimum) because the precautionary premium was built in.

Non-adjusters maintain their existing prices without the variance premium. Their gaps reflect past pricing decisions without current precautionary motives.

This divergence in starting positions drives the differential sensitivity in Theorem 1. Adjusters at elevated reset prices face symmetric exposure to future shocks—they can be pushed either up or down from their precautionary position. Non-adjusters with moderate existing gaps face asymmetric exposure—they are closer to one adjustment boundary than the other.

The curse of flexibility emerges from how different firms respond to this differential sensitivity. For flexible firms with low menu costs, the prospect of resetting to a high precautionary price that will likely become suboptimal makes adjustment less attractive. They recognise that adjusting today means accepting a position that exposes them to symmetric future shocks (higher $\partial P^a / \partial \sigma$ for $\sigma > \sigma^*$). Since they can adjust cheaply later if needed, they prefer to wait and avoid the precautionary premium altogether.

Sticky firms face the opposite calculation. Their high menu costs mean they cannot afford to wait—when their price gap becomes large, they must adjust and accept the precautionary premium despite knowing it will likely lead to future suboptimality. They cannot exploit the

differential exposure effect because the cost of maintaining current suboptimal prices exceeds the benefit of avoiding precautionary pricing.

This mechanism generates a striking selection effect: flexible firms avoid adjustment precisely to escape setting high precautionary prices. While the variance premium predicts higher reset prices conditional on adjustment, flexible firms opt for inaction instead. In contrast, sticky firms forced to adjust by large price gaps do set precautionary prices consistent with the variance premium. This selection effect—where firms most capable of adjusting choose not to, while constrained firms must accept precautionary pricing—is a direct consequence of the curse of flexibility.

Three points should be made. First, my framework explains why flexible firms reduce adjustment more than sticky firms—the variance premium in reset prices creates a stronger differential exposure effect for low-menu-cost firms who can afford to wait. Second, the non-monotonic nature of this effect is crucial: the precautionary premium grows with uncertainty, but only above the critical threshold σ^* does it dominate flexible firms’ decision calculus, triggering the curse. Below σ^* , standard menu cost intuition prevails. Third, this mechanism suggests a role for central bank communication: by reducing uncertainty about future inflation—lowering $\text{Var}_t(\pi_{t+1})$ —policymakers could alleviate the curse and restore flexible firms’ responsiveness during crises.

Next, I verify these theoretical findings numerically.

7.6. Numerical Verification of the differential sensitivity mechanism

To verify the theoretical mechanisms and ensure consistency with the empirical findings, I numerically evaluate the key relationships derived in the proof. While the stylised model is too simple to perfectly match all features of the data, it provides valuable insights into how the differential sensitivity mechanism generates the curse of flexibility.

For the menu cost calibration, I set $\psi = 0.2\%$ of revenue per adjustment for flexible firms. This choice balances several considerations. First, direct physical menu cost estimates range from 0.7-1.0% of annual revenues (see Zbaracki et al. (2004) and Levy et al. (1997)), but these are annual figures across all price changes. For a flexible firm adjusting quarterly, 0.2% per adjustment implies 0.8% annual costs, well within these estimates. Second, with profit function curvature $\kappa = 1.2$, this menu cost generates an adjustment threshold of

$\bar{x} = \sqrt{2\psi/\kappa} \approx 5.8\%$, which aligns closely with the 5-6% secondary mode observed in the distribution of price changes for high-frequency adjusters in the data. This secondary mode likely represents economically meaningful adjustments, as the 2% primary mode may reflect measurement error, rounding, or very small tactical adjustments. Third, this calibration ensures the model exhibits the curse of flexibility across a wide range of uncertainty levels, consistent with the empirical findings.

Figure 3 illustrates how the differential sensitivity mechanism generates the curse of flexibility across different uncertainty levels. I examine uncertainty levels $\sigma \in [0.02, 0.12]$, where σ represents the standard deviation of price shocks. This range spans from low volatility environments (2% shocks) to extreme volatility (12% shocks), capturing both tranquil periods and crisis-level uncertainty. and use a quarterly discount factor $\beta = 0.99$.

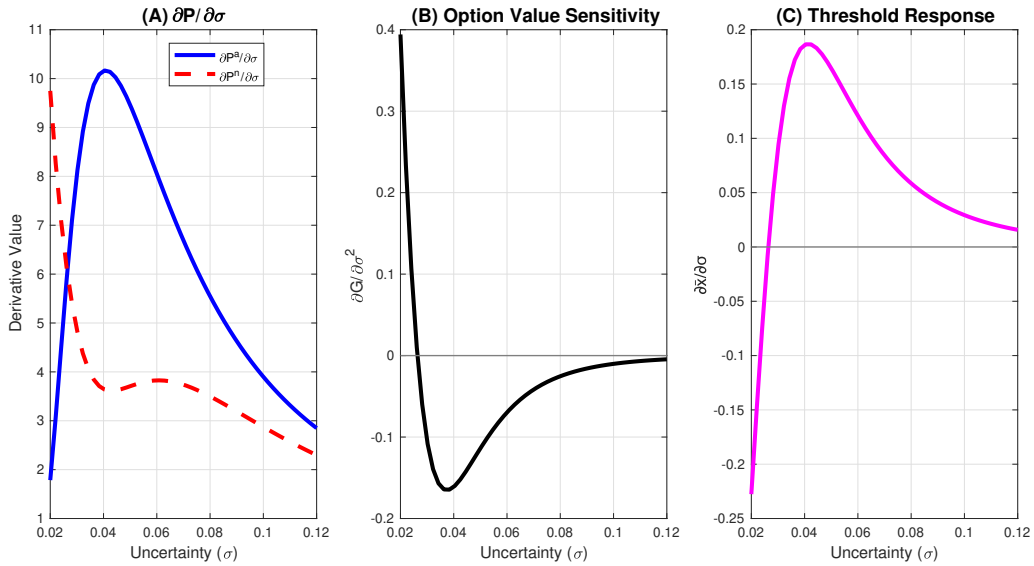


Figure 3: Mathematical Proof of the Curse of Flexibility

Notes: This figure illustrates the differential sensitivity mechanism. Panel A shows that firms which reset their prices become more exposed to future uncertainty than firms that wait. This differential exposure is strongest at moderate uncertainty levels. Panel B demonstrates that this exposure difference creates an differential exposure effect ($\partial G / \partial \sigma^2 < 0$)—adjustment becomes less attractive when uncertainty is high. Panel C shows the behavioural response: firms tolerate larger price gaps before adjusting, reducing their adjustment frequency. The curse of flexibility is strongest at moderate uncertainty ($\sigma \approx 0.04$) and persists but weakens at extreme levels. The calibration assumes flexible firms with low menu costs ($\psi = 0.2\%$), quarterly discount factor $\beta = 0.99$, and standard profit curvature.

In Panel A, the blue solid line shows $\partial P^a / \partial \sigma$, measuring how adjusters' probability of needing to adjust again responds to uncertainty changes. The red dashed line shows $\partial P^n / \partial \sigma$,

capturing the same sensitivity for non-adjusters. At very low uncertainty ($\sigma < 0.0264$), non-adjusters are more sensitive to uncertainty shocks than adjusters, consistent with standard menu cost models. This is because non-adjusters are typically closer to their adjustment thresholds due to accumulated price drift, while recent adjusters have reset to optimal prices and enjoy a buffer before needing to adjust again. The no curse result is in line with the identification strategy using calm-period flexibility.

However, this relationship reverses as uncertainty rises. Consistent with Theorem 1, we observe that the curse emerges when uncertainty increases beyond the $\sigma^* = 0.0264$ threshold. This validates the identification strategy: firms classified as flexible during tranquil periods (when no curse exists) subsequently exhibit strong curse effects when uncertainty rises, confirming these are inherent firm characteristics rather than endogenous responses. Crucially, as uncertainty increases, $\partial P^a / \partial \sigma > \partial P^n / \partial \sigma$ —the opposite of standard model predictions—confirming that adjusters are indeed more vulnerable to uncertainty than non-adjusters. While the curse is strongest at relatively moderate uncertainty levels ($\sigma = 0.04$), it remains remarkably persistent even as uncertainty increases to extreme levels. Note that $\sigma = 0.04$ is the point where the sensitivity difference peaks, not where it crosses. Consistent with the empirical findings, the curse weakens with increasing uncertainty but continues to exist even at $\sigma = 0.12$ —representing high level uncertainty.

Panel B demonstrates the differential exposure effect that emerges from this differential sensitivity. The black line shows $\partial G / \partial \sigma^2$, which measures how uncertainty affects the attractiveness of adjustment relative to waiting. The consistently negative values confirm the theoretical prediction: higher uncertainty makes adjustment less attractive because it increases the differential exposure effect. The effect reaches its minimum (most negative) around $\sigma = 0.04$, precisely where the differential sensitivity in Panel A is strongest. As uncertainty approaches extreme levels, the differential exposure effect diminishes, approaching zero as the benefits of waiting are overwhelmed by the costs of large price deviations.

Panel C presents the final result of the differential sensitivity mechanism: how firms adjust their tolerance for price gaps in response to uncertainty. The magenta line shows $\partial \bar{x} / \partial \sigma$, representing the response of adjustment thresholds to uncertainty changes. Firms raise their thresholds when facing higher uncertainty, meaning they tolerate larger price gaps before

adjusting. This threshold response exhibits an inverted-U shape, peaking around $\sigma = 0.04$, then gradually declining as uncertainty increases further. This pattern directly translates into reduced adjustment frequencies for flexible firms during uncertain times.

7.7. General Equilibrium

7.7.1. Market Clearing and Consistency Conditions

In equilibrium, three conditions must hold simultaneously. First, the labour market clears:

$$\int_0^1 n_{it} di = N_t \quad (47)$$

where N_t is aggregate labor supply. Second, the goods market clears for each variety:

$$y_{it} = c_{it} = \left(\frac{P_{it}}{P_t} \right)^{-\epsilon} C_t \quad (48)$$

Third, aggregate consistency requires:

$$Y_t = C_t \quad (49)$$

7.7.2. Distribution Dynamics

Firms are distributed across menu costs $\psi \sim F(\psi)$ and price gaps x . The adjustment probability $\lambda(\psi, \sigma_t^2)$ satisfies the flexibility curse:

$$\frac{\partial \lambda}{\partial \sigma_t^2} = \begin{cases} \approx 0 & \text{for high } \psi \text{ (sticky firms)} \\ < 0 & \text{for low } \psi \text{ (flexible firms)} \end{cases} \quad (50)$$

The joint distribution $\Gamma_t(\psi, x)$ evolves as:

$$\Gamma_{t+1}(\psi, x) = \int \Gamma_t(\psi, x') [\lambda_t \delta_0(x) + (1 - \lambda_t) \delta(x' - x - \pi_{t+1})] dx' \quad (51)$$

where δ_0 is a Dirac delta at the optimal reset price (normalized to 0) and π_{t+1} is inflation. This captures the discrete nature of adjustment: firms that adjust reset their price gaps to zero, while firms that maintain their prices see their relative gaps drift by the inflation amount π_{t+1} .

The model deliberately simplifies by having all adjusting firms reset to the same optimal price, as shown in the distribution dynamics where δ_0 represents a common reset point. This allows me to isolate how menu cost heterogeneity affects adjustment frequencies

rather than reset prices. The key insight—that flexible firms reduce adjustment under uncertainty—operates through the extensive margin (whether to adjust) rather than the intensive margin (reset price choice). While a richer model could allow menu-cost-dependent reset prices, this would complicate the analysis without changing the fundamental mechanism: flexible firms avoid adjustment to escape symmetric exposure to future shocks.

7.7.3. *Heterogeneous Responses and Aggregate Outcomes*

The aggregate adjustment rate Λ_t reflects type-dependent responses to uncertainty:

$$\Lambda_t = \int \lambda(\psi, \sigma_t^2) d\Gamma_t(\psi, x) \quad (52)$$

The flexibility curve creates non-monotonicity in the aggregate response:

$$\frac{\partial \Lambda_t}{\partial \sigma_t^2} = \int \underbrace{\frac{\partial \lambda(\psi, \sigma_t^2)}{\partial \sigma_t^2}}_{\substack{+ \text{ if } \psi \text{ high} \\ - \text{ if } \psi \text{ low}}} d\Gamma_t(\psi, x) \quad (53)$$

While standard models universally predict $\partial \Lambda_t / \partial \sigma_t^2 > 0$, my model permits $\partial \Lambda_t / \partial \sigma_t^2 < 0$ when flexible firms dominate. The sign depends on Γ_t 's mass allocation across ψ types.

8. Conclusions

I document a new empirical regularity—the curse of flexibility—whereby firms with the lowest menu costs reduce their price adjustment most when uncertainty rises. Using UK micro-price data spanning 2003-2021, I show that flexible firms decrease their adjustment probability by 1.08 percentage points more than sticky firms per standard deviation of uncertainty. While these magnitudes may appear modest, they represent a complete reversal of theoretical predictions: finding any negative differential challenges fundamental menu cost logic. The Brexit referendum provides particularly clean identification, with the differential strengthening to 2.0 percentage points when uncertainty and skewness are orthogonal. To put these numbers in perspective, uncertainty generates a larger differential response than monetary policy shocks (0.91 percentage points)—suggesting that uncertainty can dominate traditional policy channels in determining firm behaviour.

The theoretical explanation centres on differential exposure within uncertain environments. When uncertainty exceeds a critical threshold, adjusters and non-adjusters face

fundamentally different risk profiles. Firms that reset to optimal prices face symmetric exposure—equally vulnerable to positive and negative shocks around their reset position. Non-adjusters with existing price gaps face asymmetric exposure—having already drifted toward one adjustment boundary, they’re partially protected from shocks pushing them further in that direction. This asymmetry makes waiting valuable for flexible firms: unlike sticky firms constrained by high menu costs, they can maintain their protected position until conditions clarify.

These findings fundamentally challenge our understanding of price rigidity. Standard models attribute stickiness to mechanical constraints—high menu costs preventing desired adjustments. The curse reveals a distinct phenomenon: strategic rigidity, where the most flexible firms deliberately choose inaction to preserve asymmetric positions during uncertain times. This distinction has profound policy implications. Policies aimed at reducing menu costs (digitalisation, deregulation) may paradoxically reduce price flexibility during crises by creating more firms vulnerable to the curse. When uncertainty is high, having more flexible firms doesn’t guarantee more responsive prices—it may guarantee more strategic waiting.

Finally, the differential exposure mechanism likely extends beyond pricing to any decision where adjustment costs vary across agents and uncertainty affects the value of waiting—labor contracts, capital investment, inventory management, and financial portfolio rebalancing. I leave these extensions to future research.

Appendix A. Detailed Derivation of Value Function Approximation

I derive the approximation:

$$V(\hat{p}, \Omega) \approx V^*(\Omega) - \frac{1}{2}\kappa(\Omega)(\hat{p} - \hat{p}^*(\Omega))^2 \quad (\text{A.1})$$

The profit function $\Pi(\hat{p}, \Omega)$ is twice continuously differentiable and strictly concave in \hat{p} , implying $\Pi_{\hat{p}\hat{p}}(\hat{p}^*, \Omega) < 0$.

Step 1: Start with the Bellman Equation

The value function for a firm with menu cost ψ and price \hat{p} satisfies:

$$V(\hat{p}, \Omega, \psi) = \max\{V^n(\hat{p}, \Omega, \psi), V^a(\Omega, \psi) - \psi w\} \quad (\text{A.2})$$

where

$$V^n(\hat{p}, \Omega, \psi) = \Pi(\hat{p}, \Omega) + \beta\mathbb{E}[V(\hat{p} - \pi', \Omega', \psi)] \quad (\text{A.3})$$

$$V^a(\Omega, \psi) = \max_{\hat{p}'}\{\Pi(\hat{p}', \Omega) + \beta\mathbb{E}[V(\hat{p}' - \pi', \Omega', \psi)]\} \quad (\text{A.4})$$

and π' denotes next period's inflation rate.

Step 2: Focus on the Non-Adjustment Region

For prices sufficiently close to \hat{p}^* , the firm optimally chooses not to adjust. In this region:

$$V(\hat{p}, \Omega, \psi) = \Pi(\hat{p}, \Omega) + \beta\mathbb{E}[V(\hat{p} - \pi', \Omega', \psi)] \quad (\text{A.5})$$

Step 3: Define the Frictionless Value Function

Let $V^*(\Omega, \psi)$ be the value when the firm can freely set its price without paying menu costs:

$$V^*(\Omega, \psi) = \max_{\hat{p}}\{\Pi(\hat{p}, \Omega) + \beta\mathbb{E}[V(\hat{p} - \pi', \Omega', \psi)]\} \quad (\text{A.6})$$

The optimal price $\hat{p}^*(\Omega)$ satisfies the first-order condition:

$$\Pi_{\hat{p}}(\hat{p}^*, \Omega) + \beta\mathbb{E}[V_{\hat{p}}(\hat{p}^* - \pi', \Omega', \psi)] = 0 \quad (\text{A.7})$$

Note that under the stationarity assumptions and with small inflation, $\hat{p}^*(\Omega)$ approximately equals the static profit-maximizing price $\arg \max_{\hat{p}} \Pi(\hat{p}, \Omega)$, which is the ‘flexible price’ that firms track in menu cost models.

Step 4: Taylor Expand V Around the Optimum

For \hat{p} near \hat{p}^* , we can expand V :

$$V(\hat{p}, \Omega, \psi) = V(\hat{p}^*, \Omega, \psi) + V_{\hat{p}}(\hat{p}^*, \Omega, \psi)(\hat{p} - \hat{p}^*) \quad (\text{A.8})$$

$$+ \frac{1}{2} V_{\hat{p}\hat{p}}(\hat{p}^*, \Omega, \psi)(\hat{p} - \hat{p}^*)^2 + O((\hat{p} - \hat{p}^*)^3) \quad (\text{A.9})$$

Step 5: Apply the Envelope Theorem

At the optimum \hat{p}^* , the envelope theorem implies:

$$V_{\hat{p}}(\hat{p}^*, \Omega, \psi) = 0 \quad (\text{A.10})$$

This holds because \hat{p}^* maximizes V in the frictionless case.

Step 6: Derive the Second Derivative

Differentiate the Bellman equation (in the non-adjustment region) with respect to \hat{p} :

$$V_{\hat{p}}(\hat{p}, \Omega, \psi) = \Pi_{\hat{p}}(\hat{p}, \Omega) + \beta \mathbb{E}[V_{\hat{p}}(\hat{p} - \pi', \Omega', \psi)] \quad (\text{A.11})$$

Differentiate again:

$$V_{\hat{p}\hat{p}}(\hat{p}, \Omega, \psi) = \Pi_{\hat{p}\hat{p}}(\hat{p}, \Omega) + \beta \mathbb{E}[V_{\hat{p}\hat{p}}(\hat{p} - \pi', \Omega', \psi)] \quad (\text{A.12})$$

Step 7: Evaluate at the Optimum Under Approximate Stationarity

For small inflation rates and slowly evolving aggregate conditions, the value function curvature satisfies:

$$\mathbb{E}[V_{\hat{p}\hat{p}}(\hat{p}^* - \pi', \Omega', \psi)] \approx V_{\hat{p}\hat{p}}(\hat{p}^*, \Omega, \psi) \quad (\text{A.13})$$

This holds when: (i) inflation π' is small, (ii) the optimal price $\hat{p}^*(\Omega')$ evolves slowly, and (iii) the profit function curvature is stable across states.

Under this assumption:

$$V_{\hat{p}\hat{p}}(\hat{p}^*, \Omega, \psi) = \Pi_{\hat{p}\hat{p}}(\hat{p}^*, \Omega) + \beta V_{\hat{p}\hat{p}}(\hat{p}^*, \Omega, \psi) \quad (\text{A.14})$$

Solving for $V_{\hat{p}\hat{p}}$:

$$V_{\hat{p}\hat{p}}(\hat{p}^*, \Omega, \psi) = \frac{\Pi_{\hat{p}\hat{p}}(\hat{p}^*, \Omega)}{1 - \beta} \quad (\text{A.15})$$

Step 8: The High Discount Factor Approximation

Define $\kappa(\Omega) \equiv -\Pi_{\hat{p}\hat{p}}(\hat{p}^*, \Omega) > 0$ as the curvature of the profit function. Then:

$$V_{\hat{p}\hat{p}}(\hat{p}^*, \Omega, \psi) = \frac{-\kappa(\Omega)}{1 - \beta} \quad (\text{A.16})$$

For typical calibrations with quarterly discount factors $\beta \approx 0.99$:

- $1/(1 - \beta) \approx 100$ for quarterly data
- The value function is approximately 100 times more curved than the profit function
- However, for analysing adjustment decisions over short horizons (1-4 periods), the effective curvature is closer to $\kappa(\Omega)$

Since adjustment decisions typically involve horizons of 1-4 periods, I follow the menu cost literature and use:

$$V_{\hat{p}\hat{p}}(\hat{p}^*, \Omega, \psi) \approx -\kappa(\Omega) \quad (\text{A.17})$$

Step 9: Final Approximation

Substituting back into the Taylor expansion:

$$V(\hat{p}, \Omega, \psi) \approx V(\hat{p}^*, \Omega, \psi) + 0 \cdot (\hat{p} - \hat{p}^*) - \frac{1}{2}\kappa(\Omega)(\hat{p} - \hat{p}^*)^2 \quad (\text{A.18})$$

Since $V(\hat{p}^*, \Omega, \psi) = V^*(\Omega, \psi)$ by definition:

$$V(\hat{p}, \Omega, \psi) \approx V^*(\Omega, \psi) - \frac{1}{2}\kappa(\Omega)(\hat{p} - \hat{p}^*)^2 \quad (\text{A.19})$$

Appendix B. Derivation of Optimal Reset Price

Appendix B.1. Setup and First-Order Condition

The firm choosing reset price \hat{p}_t^R solves:

$$\max_{\hat{p}_t^R} \{ \Pi(\hat{p}_t^R, \Omega_t) + \beta \mathbb{E}_t[V(\hat{p}_t^R - \pi_{t+1}, \Omega_{t+1})] \} \quad (\text{B.1})$$

The first-order condition is:

$$\Pi_{\hat{p}}(\hat{p}_t^R, \Omega_t) + \beta \mathbb{E}_t[V_{\hat{p}}(\hat{p}_t^R - \pi_{t+1}, \Omega_{t+1})] = 0 \quad (\text{B.2})$$

Appendix B.2. Higher-Order Profit Function Expansion

While Appendix A derived the quadratic approximation for the value function:

$$V(\hat{p}, \Omega) \approx V^*(\Omega) - \frac{1}{2} \kappa_t (\hat{p} - \hat{p}^*)^2 \quad (\text{B.3})$$

this yields $V_{\hat{p}} \approx -\kappa_t (\hat{p} - \hat{p}^*)$. However, to capture second-order reset price effects from inflation uncertainty, I must retain higher-order terms in the **profit function** expansion.

Key Insight: The reset price depends on:

$$\Pi_{\hat{p}}(\hat{p}_t^R, \Omega_t) + \beta \mathbb{E}_t[V_{\hat{p}}(\hat{p}_t^R - \pi_{t+1}, \Omega_{t+1})] = 0 \quad (\text{B.4})$$

where:

1. $\Pi_{\hat{p}}$ requires a cubic expansion for precision
2. $V_{\hat{p}}$ can remain linear because its curvature is already captured via $\kappa_t = -\Pi_{\hat{p}\hat{p}}$

Implementation: Define $x_t \equiv \hat{p}_t^R - \hat{p}_t^*$. Since $\Pi_{\hat{p}}(\hat{p}_t^*, \Omega_t) = 0$ by optimality:

1. Expand profits to third order:

$$\Pi(\hat{p}_t^R, \Omega_t) \approx \Pi(\hat{p}_t^*, \Omega_t) - \frac{\kappa_t}{2} x_t^2 - \frac{\kappa'_t}{6} x_t^3 \quad (\text{B.5})$$

$$\Rightarrow \Pi_{\hat{p}}(\hat{p}_t^R, \Omega_t) \approx -\kappa_t x_t - \frac{\kappa'_t}{2} x_t^2 \quad (\text{B.6})$$

where $\kappa_t \equiv -\Pi_{\hat{p}\hat{p}}(\hat{p}_t^*, \Omega_t) > 0$ and $\kappa'_t \equiv -\Pi_{\hat{p}\hat{p}\hat{p}}(\hat{p}_t^*, \Omega_t)$.

2. Keep value function derivative linear:⁴

$$V_{\hat{p}}(\hat{p}_t^R - \pi_{t+1}, \Omega_{t+1}) \approx -\kappa_{t+1}(\hat{p}_t^R - \pi_{t+1} - \hat{p}_{t+1}^*) \quad (\text{B.7})$$

Appendix B.3. Expected Future Marginal Value

The future price gap is:

$$\hat{p}_t^R - \pi_{t+1} - \hat{p}_{t+1}^* = x_t - \pi_{t+1} - (\hat{p}_{t+1}^* - \hat{p}_t^*) \quad (\text{B.8})$$

⁴While one could derive $V(\hat{p})$ to third order, this would complicate Appendix A without new economic insights. My approach isolates where higher-order terms matter most.

I approximate the transition of the optimal price as:

$$\hat{p}_{t+1}^* - \hat{p}_t^* \approx \rho \pi_{t+1} \quad (\text{B.9})$$

where $\rho \in [0, 1]$ captures incomplete pass-through due to strategic complementarities. Following the literature (e.g., Wang and Werning (2022), Alvarez et al. (2023) and Gagliardone et al. (2025)), I can treat ρ as a constant parameter. Alternatively, ρ can be endogenized through the signal extraction mechanism proposed in Kara (2025), where:

$$\rho(\pi_t) = \frac{\pi_t^2}{\pi_t^2 + \sigma_{idio}^2}$$

generates regime-dependent strategic behavior, suggested by the UK micro data, with substitutability below 2% inflation and complementarity above.

Therefore:

$$\hat{p}_t^R - \pi_{t+1} - \hat{p}_{t+1}^* = x_t - (1 + \rho)\pi_{t+1} \equiv z_{t+1} \quad (\text{B.10})$$

The expected future marginal value becomes:

$$\mathbb{E}_t[V_{\hat{p}}(\hat{p}_t^R - \pi_{t+1}, \Omega_{t+1})] \approx -\mathbb{E}_t[\kappa_{t+1} z_{t+1}] \quad (\text{B.11})$$

Appendix B.4. Key Approximations

Table B.7: Summary of Approximations

Term	Approximation	Justification
Marginal profit	$\Pi_{\hat{p}} \approx -\kappa_t x_t - \frac{\kappa_t'}{2} x_t^2$	Cubic expansion
Value derivative	$V_{\hat{p}} \approx -\kappa_t (\hat{p} - \hat{p}_t^*)$	Quadratic value
Optimal price transition	$\hat{p}_{t+1}^* - \hat{p}_t^* \approx \rho \pi_{t+1}$	Linear pass-through
Profit curvature	$\kappa_{t+1} \approx \kappa_t$	Slow evolution

Under these approximations:

$$\mathbb{E}_t[V_{\hat{p}}(\hat{p}_t^R - \pi_{t+1}, \Omega_{t+1})] \approx -\kappa_t \mathbb{E}_t[z_{t+1}] \quad (\text{B.12})$$

Appendix B.5. Computing Moments

I need the first two moments of $z_{t+1} = x_t - (1 + \rho)\pi_{t+1}$:

$$\mathbb{E}_t[z_{t+1}] = x_t - (1 + \rho)\mathbb{E}_t[\pi_{t+1}] \quad (\text{B.13})$$

$$\mathbb{E}_t[z_{t+1}^2] = \mathbb{E}_t[(x_t - (1 + \rho)\pi_{t+1})^2] \quad (\text{B.14})$$

$$= x_t^2 - 2x_t(1 + \rho)\mathbb{E}_t[\pi_{t+1}] + (1 + \rho)^2\mathbb{E}_t[\pi_{t+1}^2] \quad (\text{B.15})$$

$$= x_t^2 - 2x_t(1 + \rho)\mathbb{E}_t[\pi_{t+1}] + (1 + \rho)^2[\text{Var}_t(\pi_{t+1}) + \mathbb{E}_t[\pi_{t+1}]^2] \quad (\text{B.16})$$

Appendix B.6. Substituting into the FOC

The first-order condition becomes:

$$0 = \Pi_{\hat{p}}(\hat{p}_t^R, \Omega_t) + \beta\mathbb{E}_t[V_{\hat{p}}(\hat{p}_t^R - \pi_{t+1}, \Omega_{t+1})] \quad (\text{B.17})$$

$$= -\kappa_t x_t - \frac{\kappa'_t}{2} x_t^2 - \beta\kappa_t[x_t - (1 + \rho)\mathbb{E}_t[\pi_{t+1}]] \quad (\text{B.18})$$

Rearranging:

$$0 = -(1 + \beta)\kappa_t x_t + \beta\kappa_t(1 + \rho)\mathbb{E}_t[\pi_{t+1}] - \frac{\kappa'_t}{2} x_t^2 \quad (\text{B.19})$$

Appendix B.7. Solving Order by Order

First-order solution: Ignoring the x_t^2 term initially:

$$-(1 + \beta)\kappa_t x_t + \beta\kappa_t(1 + \rho)\mathbb{E}_t[\pi_{t+1}] = 0 \quad (\text{B.20})$$

This yields:

$$x_t^{(1)} = \frac{\beta(1 + \rho)}{1 + \beta}\mathbb{E}_t[\pi_{t+1}] \equiv \alpha_1\mathbb{E}_t[\pi_{t+1}] \quad (\text{B.21})$$

Second-order correction: To capture variance effects, I expand the expectation term more carefully. Note that:

$$\mathbb{E}_t[z_{t+1}^2] = \mathbb{E}_t[(x_t - (1 + \rho)\pi_{t+1})^2] \quad (\text{B.22})$$

$$= x_t^2 - 2x_t(1 + \rho)\mathbb{E}_t[\pi_{t+1}] + (1 + \rho)^2\text{Var}_t(\pi_{t+1}) + (1 + \rho)^2\mathbb{E}_t[\pi_{t+1}]^2 \quad (\text{B.23})$$

When I substitute $x_t = x_t^{(1)} + x_t^{(2)}$ and account for the interaction between the quadratic profit term and inflation variance, I obtain:

$$x_t^{(2)} = \frac{\beta(1+\rho)^2}{2(1+\beta)} \frac{\kappa'_t}{\kappa_t} \text{Var}_t(\pi_{t+1}) \quad (\text{B.24})$$

Using the relationship $\kappa'_t/\kappa_t = \Pi_{\hat{p}\hat{p}}/\Pi_{\hat{p}} = -(\epsilon - 2)$ from the profit function analysis:

$$x_t^{(2)} = \frac{\beta(\epsilon - 2)(1+\rho)^2}{2(1+\beta)} \text{Var}_t(\pi_{t+1}) \equiv \alpha_2 \text{Var}_t(\pi_{t+1}) \quad (\text{B.25})$$

Appendix B.8. Final Result

Combining the first and second-order terms:

$$\underbrace{\hat{p}_t^R - \hat{p}_t^*}_{\text{Price gap}} = \underbrace{\frac{\beta(1+\rho)}{1+\beta} \mathbb{E}_t[\pi_{t+1}]}_{\text{Inflation expectation}} + \underbrace{\frac{\beta(\epsilon - 2)(1+\rho)^2}{2(1+\beta)} \text{Var}_t(\pi_{t+1})}_{\text{Uncertainty premium}} \quad (\text{B.26})$$

The optimal reset price has two components:

1. **First-order term:** Forward-looking adjustment for expected inflation erosion
2. **Second-order term:** Precautionary pricing due to uncertainty (positive when $\epsilon > 2$)

For $\epsilon = 2$ (log utility), $\kappa'_t = 0$ and I recover Golosov-Lucas's linear approximation.

Appendix C: Detailed Proof of $\partial G/\partial \sigma^2 < 0$.

Foundational Setup Given:

- Inflation shock: $\pi' \sim N(0, \sigma^2)$ (normally distributed with mean 0, variance σ^2)
- Adjustment threshold: \bar{x} (firm adjusts when price gap exceeds this)
- Standard normal CDF: $\Phi(z) = \Pr(Z \leq z)$ where $Z \sim N(0, 1)$

Key relationship: If $X \sim N(\mu, \sigma^2)$, then $(X - \mu)/\sigma \sim N(0, 1)$

Case 1: Adjuster Probability $P^a(\sigma) = \Pr(|\pi'| > \bar{x})$ **What I'm computing:** A firm just adjusted (gap = 0). Next period, inflation π' hits. What's the probability the new gap $|\pi'|$ exceeds the threshold \bar{x} ?

Step 1: Break down the absolute value $|\pi'| > \bar{x}$ occurs when:

- $\pi' > \bar{x}$ (positive shock large enough), OR
- $\pi' < -\bar{x}$ (negative shock large enough)

So: $\Pr(|\pi'| > \bar{x}) = \Pr(\pi' > \bar{x}) + \Pr(\pi' < -\bar{x})$

Step 2: Standardize $\pi' \sim N(0, \sigma^2)$ Define the standard normal variable: $Z = \pi'/\sigma \sim N(0, 1)$

Then: $\pi' = \sigma Z$

Step 3: Convert the inequalities

- $\pi' > \bar{x}$ becomes $\sigma Z > \bar{x}$, so $Z > \bar{x}/\sigma$
- $\pi' < -\bar{x}$ becomes $\sigma Z < -\bar{x}$, so $Z < -\bar{x}/\sigma$

Step 4: Use standard normal probabilities For the first term:

$$\Pr(\pi' > \bar{x}) = \Pr(Z > \bar{x}/\sigma) = 1 - \Pr(Z \leq \bar{x}/\sigma) = 1 - \Phi(\bar{x}/\sigma) \quad (\text{B.27})$$

Using symmetry of standard normal: $1 - \Phi(z) = \Phi(-z)$

So: $\Pr(\pi' > \bar{x}) = \Phi(-\bar{x}/\sigma)$

For the second term:

$$\Pr(\pi' < -\bar{x}) = \Pr(Z < -\bar{x}/\sigma) = \Phi(-\bar{x}/\sigma) \quad (\text{B.28})$$

Step 5: Combine

$$\Pr(|\pi'| > \bar{x}) = \Phi(-\bar{x}/\sigma) + \Phi(-\bar{x}/\sigma) = 2\Phi(-\bar{x}/\sigma) \quad (\text{B.29})$$

This result relies crucially on symmetry. Because (1) inflation shocks are symmetric around zero ($\pi' \sim N(0, \sigma^2)$), and (2) the adjustment thresholds are symmetric ($\pm\bar{x}$), the probability of exceeding the threshold in either direction is identical. This symmetry simplifies the analysis considerably—I only need to calculate one tail probability and double it. Intuitively, adjusters face maximum exposure to uncertainty. Having just reset to the optimal price (gap = 0), they are vulnerable to shocks from both sides of the distribution. Any inflation shock—positive or negative—moves them away from their target. This bilateral exposure is why adjusters' probability depends on both tails of the distribution.

Case 2: Non-Adjuster Probability $P^n(x, \sigma) = \Pr(|x - \pi'| > \bar{x})$ **What I'm computing:** A firm didn't adjust (current gap = x). Next period, inflation π' hits, making the new gap $(x - \pi')$. What's the probability this new gap exceeds threshold \bar{x} ?

Step 1: Break down the absolute value $|x - \pi'| > \bar{x}$ occurs when:

- $(x - \pi') > \bar{x}$, OR
- $(x - \pi') < -\bar{x}$

Step 2: Solve the inequalities First inequality: $x - \pi' > \bar{x}$

$$x - \pi' > \bar{x} \tag{B.30}$$

$$-\pi' > \bar{x} - x \tag{B.31}$$

$$\pi' < x - \bar{x} \tag{B.32}$$

Second inequality: $x - \pi' < -\bar{x}$

$$x - \pi' < -\bar{x} \tag{B.33}$$

$$-\pi' < -\bar{x} - x \tag{B.34}$$

$$\pi' > \bar{x} + x \tag{B.35}$$

So: $\Pr(|x - \pi'| > \bar{x}) = \Pr(\pi' < x - \bar{x}) + \Pr(\pi' > \bar{x} + x)$

Step 3: Standardize using $Z = \pi'/\sigma$

- $\pi' < x - \bar{x}$ becomes $Z < (x - \bar{x})/\sigma$
- $\pi' > \bar{x} + x$ becomes $Z > (\bar{x} + x)/\sigma$

Step 4: Use standard normal probabilities First term:

$$\Pr(\pi' < x - \bar{x}) = \Pr(Z < (x - \bar{x})/\sigma) = \Phi\left(\frac{x - \bar{x}}{\sigma}\right) \tag{B.36}$$

Second term:

$$\Pr(\pi' > \bar{x} + x) = \Pr(Z > (\bar{x} + x)/\sigma) = 1 - \Phi\left(\frac{\bar{x} + x}{\sigma}\right) = \Phi\left(-\frac{\bar{x} + x}{\sigma}\right) \tag{B.37}$$

Step 5: Combine

$$\Pr(|x - \pi'| > \bar{x}) = \Phi\left(\frac{x - \bar{x}}{\sigma}\right) + \Phi\left(-\frac{\bar{x} + x}{\sigma}\right) \quad (\text{B.38})$$

Computing Derivatives with Respect to σ **Step 6: Differentiate the adjuster probability**

I need to compute $\frac{\partial P^a}{\partial \sigma}$ where $P^a(\sigma) = 2\Phi(-\bar{x}/\sigma)$.

Apply the chain rule: The chain rule states that for a composite function $f(g(x))$, I have $\frac{d}{dx}[f(g(x))] = f'(g(x)) \cdot g'(x)$.

In this case:

- Outer function: $f(u) = 2\Phi(u)$, so $f'(u) = 2\phi(u)$
- Inner function: $g(\sigma) = -\bar{x}/\sigma$

Therefore:

$$\frac{\partial P^a}{\partial \sigma} = \frac{\partial}{\partial \sigma} [2\Phi(-\bar{x}/\sigma)] = 2\phi(-\bar{x}/\sigma) \cdot \frac{\partial}{\partial \sigma}(-\bar{x}/\sigma) \quad (\text{B.39})$$

Compute the inner derivative: I need $\frac{\partial}{\partial \sigma}(-\bar{x}/\sigma)$. Since \bar{x} is a constant with respect to σ : $\frac{\partial}{\partial \sigma}(-\bar{x}/\sigma) = \frac{\bar{x}}{\sigma^2}$

Combine the results: Substituting back into the chain rule expression:

$$\frac{\partial P^a}{\partial \sigma} = 2\phi(-\bar{x}/\sigma) \cdot \frac{\bar{x}}{\sigma^2} \quad (\text{B.40})$$

$$= \frac{2\bar{x}}{\sigma^2} \phi(-\bar{x}/\sigma) \quad (\text{B.41})$$

Apply symmetry of the standard normal density: The standard normal probability density function has the property that $\phi(-z) = \phi(z)$ for any z (it's symmetric around zero). Therefore:

$$\phi(-\bar{x}/\sigma) = \phi(\bar{x}/\sigma) \quad (\text{B.42})$$

Final result:

$$\frac{\partial P^a}{\partial \sigma} = \frac{2\bar{x}}{\sigma^2} \phi\left(\frac{\bar{x}}{\sigma}\right) \quad (\text{B.43})$$

This derivative tells us how the adjuster's probability of needing to adjust again responds to changes in uncertainty σ .

Step 7: Differentiate the non-adjuster probability

$$\frac{\partial P^n}{\partial \sigma} = \frac{\partial}{\partial \sigma} \left[\Phi \left(\frac{x - \bar{x}}{\sigma} \right) + \Phi \left(-\frac{\bar{x} + x}{\sigma} \right) \right] \quad (\text{B.44})$$

$$= \phi \left(\frac{x - \bar{x}}{\sigma} \right) \cdot \frac{-(x - \bar{x})}{\sigma^2} + \phi \left(-\frac{\bar{x} + x}{\sigma} \right) \cdot \frac{\bar{x} + x}{\sigma^2} \quad (\text{B.45})$$

$$= \frac{\bar{x} - x}{\sigma^2} \phi \left(\frac{\bar{x} - x}{\sigma} \right) + \frac{\bar{x} + x}{\sigma^2} \phi \left(\frac{\bar{x} + x}{\sigma} \right) \quad (\text{B.46})$$

Key observation: For non-adjusters with $x \in (0, \bar{x})$, the derivative $\frac{\partial P^n}{\partial \sigma}$ is positive, meaning higher uncertainty increases their adjustment probability. However, this increase is smaller than for adjusters because:

- Non-adjusters need more extreme shocks to hit the boundaries $\pm \bar{x}$ from their starting position x
- The two terms have different magnitudes: shocks need to be size $(x - \bar{x})$ in one direction but $(x + \bar{x})$ in the other
- This asymmetry provides partial insulation compared to adjusters who face symmetric exposure

Appendix B.9. Analytical Proof: How Sensitivities Depend on σ

Appendix B.10. The Key Inequality

We need to determine when:

$$\frac{\partial P^a}{\partial \sigma} \stackrel{?}{>} \frac{\partial P^n}{\partial \sigma} \quad (\text{B.47})$$

Which is:

$$\frac{2\bar{x}}{\sigma^2} \phi \left(\frac{\bar{x}}{\sigma} \right) \stackrel{?}{>} \frac{\bar{x} - x}{\sigma^2} \phi \left(\frac{\bar{x} - x}{\sigma} \right) + \frac{\bar{x} + x}{\sigma^2} \phi \left(\frac{\bar{x} + x}{\sigma} \right) \quad (\text{B.48})$$

Throughout the analysis, I consider the economically relevant case where $x \in (0, \bar{x})$, i.e., the non-adjuster's price gap is below the adjustment threshold.

Appendix B.11. Step 1: Factor Out Common Terms

Multiply both sides by σ^2 and use $\phi(z) = \frac{1}{\sqrt{2\pi}} e^{-z^2/2}$:

$$2\bar{x} \cdot e^{-\bar{x}^2/(2\sigma^2)} \stackrel{?}{>} (\bar{x} - x) \cdot e^{-(\bar{x}-x)^2/(2\sigma^2)} + (\bar{x} + x) \cdot e^{-(\bar{x}+x)^2/(2\sigma^2)} \quad (\text{B.49})$$

Appendix B.12. Step 2: Define a Key Variable

Let $t = 1/(2\sigma^2)$. As σ varies:

- When $\sigma \rightarrow 0$: $t \rightarrow \infty$ (large t)
- When $\sigma \rightarrow \infty$: $t \rightarrow 0$ (small t)

The inequality becomes:

$$2\bar{x} \cdot e^{-\bar{x}^2 t} \stackrel{?}{>} (\bar{x} - x) \cdot e^{-(\bar{x}-x)^2 t} + (\bar{x} + x) \cdot e^{-(\bar{x}+x)^2 t} \quad (\text{B.50})$$

Appendix B.13. Step 3: Analyze the Limiting Cases

Appendix B.13.1. Case A: Small σ (Large t)

Divide both sides by $e^{-\bar{x}^2 t}$:

$$2\bar{x} \stackrel{?}{>} (\bar{x} - x) \cdot e^{[\bar{x}^2 - (\bar{x}-x)^2]t} + (\bar{x} + x) \cdot e^{[\bar{x}^2 - (\bar{x}+x)^2]t} \quad (\text{B.51})$$

Calculate the exponents:

$$\bar{x}^2 - (\bar{x} - x)^2 = \bar{x}^2 - (\bar{x}^2 - 2\bar{x}x + x^2) = 2\bar{x}x - x^2 > 0 \quad \text{for } x \in (0, \bar{x}) \quad (\text{B.52})$$

$$\bar{x}^2 - (\bar{x} + x)^2 = \bar{x}^2 - (\bar{x}^2 + 2\bar{x}x + x^2) = -2\bar{x}x - x^2 < 0 \quad (\text{B.53})$$

As $t \rightarrow \infty$:

- First RHS term: $(\bar{x} - x) \cdot e^{(2\bar{x}x - x^2)t} \rightarrow \infty$
- Second RHS term: $(\bar{x} + x) \cdot e^{-(2\bar{x}x + x^2)t} \rightarrow 0$

Result: LHS = $2\bar{x} < \infty$ = RHS

Therefore: $\frac{\partial P^a}{\partial \sigma} < \frac{\partial P^n}{\partial \sigma}$ for small σ

Appendix B.13.2. Case B: Large σ (Small t)

When σ is large, $t = \frac{1}{2\sigma^2}$ becomes small. I use the first-order Taylor expansion $e^{-at} \approx 1 - at$ for small t .

Step 1: Expand LHS

$$\text{LHS} = 2\bar{x} \cdot e^{-\bar{x}^2 t} \quad (\text{B.54})$$

$$\approx 2\bar{x}(1 - \bar{x}^2 t) \quad (\text{B.55})$$

$$= 2\bar{x} - 2\bar{x}^3 t \quad (\text{B.56})$$

Step 2: Expand RHS, first term

$$(\bar{x} - x) \cdot e^{-(\bar{x}-x)^2 t} \approx (\bar{x} - x)[1 - (\bar{x} - x)^2 t] \quad (\text{B.57})$$

$$= (\bar{x} - x) - (\bar{x} - x)^3 t \quad (\text{B.58})$$

Step 3: Expand RHS, second term

$$(\bar{x} + x) \cdot e^{-(\bar{x}+x)^2 t} \approx (\bar{x} + x)[1 - (\bar{x} + x)^2 t] \quad (\text{B.59})$$

$$= (\bar{x} + x) - (\bar{x} + x)^3 t \quad (\text{B.60})$$

Step 4: Combine RHS terms

$$\text{RHS} = [(\bar{x} - x) - (\bar{x} - x)^3 t] + [(\bar{x} + x) - (\bar{x} + x)^3 t] \quad (\text{B.61})$$

$$= [(\bar{x} - x) + (\bar{x} + x)] - [(\bar{x} - x)^3 + (\bar{x} + x)^3]t \quad (\text{B.62})$$

$$= 2\bar{x} - [(\bar{x} - x)^3 + (\bar{x} + x)^3]t \quad (\text{B.63})$$

Step 5: Expand the cubic sum

Using binomial expansion:

$$(\bar{x} - x)^3 = \bar{x}^3 - 3\bar{x}^2 x + 3\bar{x}x^2 - x^3 \quad (\text{B.64})$$

$$(\bar{x} + x)^3 = \bar{x}^3 + 3\bar{x}^2 x + 3\bar{x}x^2 + x^3 \quad (\text{B.65})$$

Adding these:

$$(\bar{x} - x)^3 + (\bar{x} + x)^3 = 2\bar{x}^3 + 6\bar{x}x^2 \quad (\text{B.66})$$

(Note: The $\pm 3\bar{x}^2 x$ and $\pm x^3$ terms cancel)

Therefore:

$$\text{RHS} \approx 2\bar{x} - (2\bar{x}^3 + 6\bar{x}x^2)t \quad (\text{B.67})$$

Step 6: Compute the difference

$$\text{LHS} - \text{RHS} = [2\bar{x} - 2\bar{x}^3t] - [2\bar{x} - (2\bar{x}^3 + 6\bar{x}x^2)t] \quad (\text{B.68})$$

$$= 2\bar{x} - 2\bar{x}^3t - 2\bar{x} + (2\bar{x}^3 + 6\bar{x}x^2)t \quad (\text{B.69})$$

$$= -2\bar{x}^3t + 2\bar{x}^3t + 6\bar{x}x^2t \quad (\text{B.70})$$

$$= 6\bar{x}x^2t > 0 \quad (\text{B.71})$$

Since $\bar{x} > 0$, $x > 0$ (by assumption), and $t > 0$, we have $\text{LHS} > \text{RHS}$.

Result: For large σ , adjusters are more sensitive than non-adjusters: $\frac{\partial P^a}{\partial \sigma} > \frac{\partial P^n}{\partial \sigma}$

Appendix B.14. Economic Interpretation

The relationship between adjuster and non-adjuster sensitivities fundamentally depends on the level of uncertainty σ . In the low uncertainty regime ($\sigma < \sigma^*$), non-adjusters are more sensitive to uncertainty changes because they start from a position closer to one boundary, making them highly responsive to small shocks. As uncertainty increases past the critical threshold σ^* , adjusters become more sensitive due to their symmetric exposure from the zero starting point. However, this advantage diminishes with further increases in uncertainty. In the extreme uncertainty limit ($\sigma \rightarrow \infty$), the sensitivity difference vanishes entirely as $\Delta(\sigma) \approx 3\bar{x}x^2/\sigma^2 \rightarrow 0$, indicating that both types become equally affected by uncertainty and the differential sensitivity mechanism ceases to operate.

References

- Alvarez, F., Le Bihan, H., Lippi, F., 2016. The real effects of monetary shocks in sticky price models: A sufficient statistic approach. *American Economic Review* 106, 2817–2851.
- Alvarez, F., Lippi, F., Souganidis, P., 2023. Price setting with strategic complementarities as a mean field game. *Econometrica* 91, 2005–2039.
- Aruoba, S.B., Oue, E., Saffie, F., Willis, J., 2025. Reviving Micro Real Rigidities: The Importance of Demand Shocks. Working Paper 32518. National Bureau of Economic Research. First version May 2024, Revised January 2025.
- Bloom, N., 2009. The impact of uncertainty shocks. *Econometrica* 77, 623–685.
- Braun, R., Miranda-Agrippino, S., Saha, T., 2025. Measuring monetary policy in the uk: The uk monetary policy event-study database. *Journal of Monetary Economics* 149, 103645.
- Calvo, G.A., 1983. Staggered prices in a utility-maximizing framework. *Journal of Monetary Economics* 12, 383–398.
- Caplin, A.S., Spulber, D.F., 1987. Menu costs and the neutrality of money. *The Quarterly Journal of Economics* 102, 703–725.
- Davies, R., 2021. Prices and inflation in a pandemic: A micro data approach. URL: <http://cep.lse.ac.uk/pubs/download/dp017.pdf>.
- Gagliardone, L., Gertler, M., Lenzu, S., Tielens, J., 2025. Anatomy of the phillips curve: Micro evidence and macro implications. *American Economic Review* Forthcoming.
- Gilchrist, S., Schoenle, R., Sim, J., Zakrajšek, E., 2017. Inflation dynamics during the financial crisis. *American Economic Review* 107, 785–823.
- Golosov, M., Lucas, R.E., 2007. Menu costs and phillips curves. *Journal of Political Economy* 115, 171–199.
- Kara, E., 2025. The natural rate of inflation. Working Paper .
- Levy, D., Bergen, M., Dutta, S., Venable, R., 1997. The magnitude of menu costs: Direct evidence from large u.s. supermarket chains. *The Quarterly Journal of Economics* 112, 791–825.
- Sheshinski, E., Weiss, Y., 1977. Inflation and costs of price adjustment. *The Review of*

Economic Studies 44, 287–303.

Tenreyro, S., Thwaites, G., 2016. Pushing on a string: Us monetary policy is less powerful in recessions. *American Economic Journal: Macroeconomics* 8, 43–74. doi:10.1257/mac.20150016.

Vavra, J., 2014. Inflation dynamics and time-varying volatility: New evidence and an ss interpretation. *The Quarterly Journal of Economics* 129, 215–258.

Wang, O., Werning, I., 2022. Dynamic oligopoly and price stickiness. *American Economic Review* 112, 2815–49. doi:10.1257/aer.20201739.

Zbaracki, M.J., Ritson, M., Levy, D., Dutta, S., Bergen, M., 2004. Managerial and customer costs of price adjustment: Direct evidence from industrial markets. *The Review of Economics and Statistics* 86, 514–533.